LINEAR ALGEBRA OVER SEMIRINGS

A thesis submitted to The University of Manchester
for the degree of Doctor of Philosophy
in the Faculty of Engineering and Physical Sciences

DAVID WILDING

School of Mathematics

2014

Contents

Sections 1 to 20 contain approximately 37,000 words.

Abstract

Linear Algebra Over Semirings was submitted by David Wilding to The University of Manchester on 10 September 2014 for the degree of Doctor of Philosophy.

Motivated by results of linear algebra over fields, rings and tropical semirings, we present a systematic way to understand the behaviour of matrices with entries in an arbitrary semiring. We focus on three closely related problems concerning the row and column spaces of matrices. This allows us to isolate and extract common properties that hold for different reasons over different semirings, yet also lets us identify which features of linear algebra are specific to particular types of semiring. For instance, the row and column spaces of a matrix over a field are isomorphic to each others' duals, as well as to each other, but over a tropical semiring only the first of these properties holds in general (this in itself is a surprising fact). Instead of being isomorphic, the row space and column space of a tropical matrix are anti-isomorphic in a certain order-theoretic and algebraic sense.

The first problem is to describe the kernels of the row and column spaces of a given matrix. These equivalence relations generalise the orthogonal complement of a set of vectors, and the nature of their equivalence classes is entirely dependent upon the kind of semiring in question. The second, Hahn-Banach type, problem is to decide which linear functionals on row and column spaces of matrices have a linear extension. If they all do, the underlying semiring is called exact, and in this case the row and column spaces of any matrix are isomorphic to each others' duals. The final problem is to explain the connection between the row space and column space of each matrix. Our notion of a conjugation on a semiring accounts for the different possibilities in a unified manner, as it guarantees the existence of bijections between row and column spaces and lets us focus on the peculiarities of those bijections.

Our main original contribution is the systematic approach described above, but along the way we establish several new results about exactness of semirings. We give sufficient conditions for a subsemiring of an exact semiring to inherit exactness, and we apply these conditions to show that exactness transfers to finite group semirings. We also show that every Boolean ring is exact. This result is interesting because it allows us to construct a ring which is exact (also known as FP-injective) but not self-injective. Finally, we consider exactness for residuated lattices, showing that every involutive residuated lattice is exact. We end by showing that the residuated lattice of subsets of a finite monoid is exact if and only if the monoid is a group.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright

Acknowledgements

The research contained in this thesis was generously supported by The University of Manchester Faculty of Engineering and Physical Sciences.

My supervisor, Mark Kambites, suggested exactness of semirings as a potential research topic long before I even began my PhD study, and I would like to express my great appreciation to him for guiding me down such an interesting path. I would also like to deeply thank Marianne Johnson, as together the pair have been a constant source of inspiration. I have always found them very keen to discuss my ideas and involve me in their research. Working with Marianne and Mark has been an enormously enjoyable experience, and this thesis would be nowhere near what it is today without their input and advice.

Another person who always enthusiastically engaged with my research is Amit Kuber. During our years together at Manchester we have had countless discussions about our work, and I will miss our mathematical joint ventures immensely. I am also very grateful to Philip Bridge, Anthony Chiu, Andrew Davies, Markus Pfeiffer, Laura Phillips, Joe Razavi, Harold Simmons, Matthew Taylor, Marcus Tressl and David Ward for all the useful and interesting discussions we have had.

The School of Mathematics at Manchester has a large number of research groups, and I have learnt a great deal as a result of my involvement in several of them. The Mathematical Logic Group, the Tropical Mathematics Reading Group and the Centre for Interdisciplinary Computational and Dynamical Analysis made me feel extremely welcome. I also owe a large part of my mathematical development to the Mathematical Foundations Group (School of Computer Science) and the North British Semigroups and Applications Network.

I would finally like to thank my family and friends. It would not have been possible to complete this thesis without their welcome distractions and encouragement.

Introduction

1 Background and motivation

A semiring is an algebraic structure in which we can add and multiply elements, where multiplication distributes over addition, but in which neither subtraction nor division are necessarily possible. These assumptions might appear prohibitively weak; what interesting facts about semirings can we hope to establish if we cannot even rely on subtraction? But, of course, we do not usually attempt to prove grand results about all structures of a particular kind (sets, groups, rings, and so on). Rather, we focus our energy on structures that are specialised enough to allow meaningful, non-trivial, things to be said about them (well-ordered sets, simple groups, Noetherian rings). With this precedent in mind, the definition of a semiring is entirely fit for purpose. It is weak enough to encompass an extraordinary variety of mathematical objects, yet just strong enough to provide a general framework for matrices and linear algebra that does not need to be rebuilt time and time again as the class of semirings is charted.

Semirings were first explicitly defined and deemed worthy of study by Vandiver [82] in 1934, but the notion is so simple and pervasive that several authors before and since independently came very close to formulating a modern abstract definition. For instance, in 1847 Boole [11] axiomatised what we would now call an idempotent semiring (see page 15), and later Hilbert [39] and Huntington [42, 43] formulated axioms for the semiring of non-negative integers. Golan [34] gives a brief history of semiring theory, along with a very useful guide to the various alternative names under which semirings continue to appear. For the brave reader, Głazek [31] provides a comprehensive catalogue of the extensive and disparate literature on semirings and their applications. We cannot possibly describe all the areas of mathematics (and beyond) that have found a use for semirings, so in this section we only discuss a few instances when matrices over semirings are of particular importance.

One of the classic areas of mathematics in which semirings arise is the theory of formal languages. A (formal) language is any set of words—finite strings of symbols—taken from a fixed finite set, called the alphabet. For example, `a`, `bad` and `cddd` are words over the alphabet $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$, and as such the set {`a`, `bad`, `cddd`} is a language. The set of all languages over a fixed alphabet can be viewed as a semiring: to "add" two languages simply take their union, and to "multiply" two languages take the set of concatenations of a word from one language with a word from the other language (so that, for example, the product of $\{\mathtt{a}, \mathtt{b}\}$ and $\{\mathtt{c}, \mathtt{d}\}$ is $\{\mathtt{ac}, \mathtt{ad}, \mathtt{bc}, \mathtt{bd}\}$). Such semirings have applications in logic and theoretical computer science because they make it possible to compare the power of different methods of computation (see Sipser [80]).

Since a language is just an arbitrary set of words over a fixed alphabet, individual languages tend not to be as interesting as whole sets of languages that have some feature in common. One such set of languages is known as the regular languages. As a subset of the semiring of all languages, the set of regular languages is closed under addition and multiplication, so is a semiring in its own right. In fact, it is defined to be the smallest such semiring that (to continue our example) contains the special languages $\emptyset$, $\{\mathtt{a}\}$, $\{\mathtt{b}\}$, $\{\mathtt{c}\}$ and $\{\mathtt{d}\}$, and that is also closed under the unary operation which takes a language to the union of its nth powers for all integers $n \geq 0$. This operation is called the Kleene star of a language and it takes the singleton $\{\mathtt{a}\}$ to the infinite language $\{\varepsilon, \mathtt{a}, \mathtt{aa}, \mathtt{aaa}, \dots\}$, where ε denotes the unique word comprising no symbols. It is not at all obvious, but as we now explain, regular languages are intimately linked with matrices over a certain semiring.

A finite automaton is an abstract computing machine that takes as input a word over a fixed alphabet and returns as output either a 1 ("true") or a 0 ("false"). The automaton decides the fate of a word by reading it symbol by symbol and, at each step, modifying its internal state according to predefined rules. Some of the (finitely many) states the automaton can be in are designated accept states, so if the automaton finds itself in one of these states upon reaching the end of the input word it outputs a 1. Otherwise it outputs a 0. The set of words for which an automaton outputs a 1 is called the language accepted by the automaton, and by a celebrated result of Kleene [53], the regular languages are precisely the languages accepted by automata. Moreover, it turns out that the operation of an automaton can be simulated by repeatedly multiplying the Boolean matrices that encode its

rules (see Kuich and Salomaa [55]). This means that the study of regular languages essentially boils down to considering finite collections of matrices over the Boolean semiring $\{0, 1\}$, where addition is maximum ("or") and multiplication is minimum ("and").

Arguably the most obvious and well-known application of matrices with entries in a semiring is in linear algebra over a field—usually the field of real or complex numbers—where a matrix corresponds to a system of linear equations that needs to be solved. Since a field is a highly specialised type of semiring, matrices over fields have many useful and interesting properties that are not necessarily available in the case of an arbitrary semiring. We now briefly recall some of these properties (see Roman [72] for the details).

Let A denote a matrix with entries in a field, and suppose that A has $m \geq 1$ rows and $n \geq 1$ columns. Suppose further that A has rank $r \geq 1$. This means that the vector space of all linear combinations of the rows of A, i.e., the row space of A, has dimension r. The orthogonal complement of the row space of A, i.e., the null space of A, then has dimension $n - r$, and in turn the orthogonal complement of the null space of A has dimension $n - (n - r) = r$. In fact, this last vector space is equal to the row space of A, and thus applying the double orthogonal complement construction to the row space of a matrix over a field has no effect. Finally, the rank of A turns out to also be the dimension of the column space of A. Therefore the row space and column space of A are isomorphic vector spaces.

Another way to interpret this result is in terms of dual vector spaces. If X is a vector space over a field then the dual of X is defined to be the vector space of all linear functions from X to the field (linear functionals), and it is well-known that if X is finite dimensional then X is isomorphic to its dual—although the isomorphism depends on a choice of basis. In particular, the row space and column space of a matrix A are isomorphic to their respective duals. So, since the row space of A is isomorphic to the column space of A as well, we can treat the row and column spaces as if they were each others' duals. This property of matrices is of great interest because it also holds over certain semirings that are quite unlike fields, despite the fact that the row and column spaces of a matrix are not isomorphic in general. For instance, it follows from the work of Wang [83] that the Boolean semiring is such a semiring, and, more recently, some of the so-called tropical semirings have been shown to enjoy the same property (see page 15).

Tropical algebra is a relatively new area of mathematics which brings together ideas from algebra, order theory and discrete mathematics, and which has numerous applications in (for example) scheduling and optimisation, formal language theory, numerical analysis and dynamical systems. The primary objects of study are the tropical semirings—a family of semirings based on either the non-negative integers, the integers or the real numbers, but with unusual semiring operations.[1] Specifically, "addition" is either maximum or minimum (depending on the author's preference or the desired application) and "multiplication" is usual addition. That is, in the so-called max-plus formulation of tropical algebra the sum and product of two numbers a and b are given by $\max\{a, b\}$ and $a + b$ respectively, whereas in the min-plus formulation their sum and product are given by $\min\{a, b\}$ and $a + b$ respectively. To help avoid confusion, it has become commonplace to use special symbols for these repurposed operations: $a \oplus b$ means either $\max\{a, b\}$ or $\min\{a, b\}$, while $a \odot b$ and $a \otimes b$ both mean $a + b$. This process of replacing addition by maximum, say, and multiplication by addition can be thought of as taking limits of logarithms (see Litvinov [61]).

Tropical algebra is a powerful tool because it allows us to analyse inherently non-linear problems in a linear, combinatorial way. The general strategy is to first transform a (classical) non-linear system into a piecewise (tropical) linear system and then use methods from tropical linear algebra to provide information about the original system. This approach has been used to significantly speed up computation of the eigenvalues of matrix polynomials (see Gaubert and Sharify [30]), and it can also be used to understand discrete event dynamical systems (see Baccelli et al. [4]).

As a direct consequence of the usefulness of tropical linear algebra, matrices over tropical semirings have been the subject of active investigation since the 1960s. The first comprehensive account of tropical matrices and their practical applications in scheduling was compiled by Cuninghame-Green [22] in 1979, and in the years that followed, the theoretical and computational aspects of tropical linear algebra were developed to the point where they could be used to attack an enormous variety of synchronisation, optimisation and network flow problems (see Heidergott et al. [37] and Butkovič [16]). Many such problems ultimately necessitate describing the eigenvectors of a tropical matrix, and it turns out that an operation analogous to the Kleene star (see above) can be used for this purpose. Specifically, if A is a tropical

[1] Sometimes also with a minimum element $-\infty$ and/or a maximum element ∞.

matrix then the matrix $A^0 \oplus A^1 \oplus A^2 \oplus A^3 \oplus \cdots$ can—provided it exists—be used to construct the tropical eigenvectors of A. The appearance of the Kleene star here is not entirely coincidental; there is a connection between tropical algebra and formal language theory.

Recall that an automaton is an abstract machine that assigns either a 0 or a 1 to each word written using symbols taken from an underlying alphabet. So, put another way, an automaton simply computes a function from the set of words to the Boolean semiring $\{0, 1\}$. This way of thinking leads naturally to the notion of a weighted automaton—a machine that computes a function from the set of words to an arbitrary semiring—and, just as with Boolean automata, a weighted automaton can be represented by matrices with entries in the chosen output semiring S. In 1961, Schützenberger [74] showed that the set of functions to S which can be computed using weighted automata coincides with the smallest semiring of functions to S that contains certain basic functions and is closed under an appropriate Kleene star operation. This semiring is now known as the semiring of rational power series with coefficients in S, and in the case S is the Boolean semiring it is essentially just the semiring of regular languages (see Berstel and Reutenauer [7]).[1] Several finiteness problems in formal language theory have been tackled by taking S to be a tropical semiring (see Simon [78] and Pin [71]).

Various attempts have been made to use ideas from combinatorics and abstract algebra to understand the behaviour of tropical matrices. For example, d'Alessandro and Pasku [23] considered permutation properties of matrix products, Gaubert and Katz [29] investigated reachability via matrix multiplication, and Simon [79] studied finiteness conditions for semigroups of matrices. Systematic analysis of the structure of multiplicative semigroups of tropical matrices began in 2010, with semigroups of 2×2 matrices the first to be considered. Izhakian and Margolis [44] discovered semigroup identities and obtained an embedding of the bicyclic monoid, while Johnson and Kambites [46] gave a description of maximal subgroups and Green's relations. Following on from these results, Hollings and Kambites [40] and Johnson and Kambites [47] have shown that Green's relations for $n \times n$ tropical matrices are very much like those for matrices over a field. Moreover, implicit in the work of Hollings and Kambites [40] is the fact that the row and column spaces of a tropical matrix are isomorphic to each others' duals, just as with matrices over a field (see above).

[1] Functions from the set of words to S are usually called (formal) power series with coefficients in S.

Recent investigations of tropical matrices have focussed on the combinatorial and geometric properties of their row and column spaces. This approach was initiated in 2004 by Develin and Sturmfels [24], who introduced tropical polytopes (i.e., row and column spaces of tropical matrices) and studied them as polyhedral complexes using techniques from combinatorial geometry. The ability to treat a tropical polytope as a geometric object has resulted in several interlocking definitions of the dimension of a tropical polytope, and consequently there are numerous non-equivalent meanings of the rank of a tropical matrix (see Develin et al. [25] and Akian et al. [1]). In this potentially confusing environment, idempotent matrices are of great value because they essentially have only one sensible notion of rank. Idempotent matrices are also important for another reason: Johnson and Kambites [48] have shown that every finite metric space is realised on the vertices of the tropical polytope associated with some idempotent matrix of a certain kind.

Tropical geometry is a lively and rapidly expanding area of mathematics, and our discussion of it barely scratches the surface. One of the most active and productive avenues of research is the application of tropical methods in algebraic geometry, whereby algebraic varieties become polyhedral complexes called tropical varieties (see Maclagan and Sturmfels [64] for a highly readable introduction). Since tropical mathematics is inherently combinatorial, many of its successful applications in algebraic geometry involve counting the number of objects of a specific type. For example, Mikhalkin [68] has shown that complex algebraic curves in the plane can be counted by first replacing them with certain piecewise linear curves called tropical curves.

Many recently discovered aspects of tropical linear algebra were either already known to hold in a more general setting, or have been swiftly generalised after the fact. The former aspects include the techniques used to solve matrix equations and inequalities, while the latter include the relationships between row spaces, column spaces and their duals. The solution of tropical matrix equations and inequalities involves well-developed ideas from residuation theory (see Cuninghame-Green [22]), and nothing more sophisticated, so this elementary aspect of tropical linear algebra is present in a much wider variety of scenarios. We discuss some of these other applications of residuation theory below. The duals of row and column spaces of tropical matrices can also be described using residuation theory, but this relies upon a special feature of the tropical semirings—namely the ability to negate elements.

The observation that negation is of fundamental importance in tropical algebra has prompted several researchers to investigate idempotent semirings which possess a negation-like operation. A semiring is called idempotent if, as with the tropical semiring, its addition operation is idempotent. That is, a semiring is idempotent if it satisfies $a + a = a$ for all elements a (note that for tropical semirings this means $a \oplus a = a$, not $a \odot a = a$). An idempotent semiring is called complete if arbitrary sums are possible. This condition forces the semiring to have a minimum element (the empty sum) and a maximum element (the sum of the whole semiring) in the order defined by $a \leq b$ if and only if $a + b = b$. For example, the Boolean semiring is a complete idempotent semiring with minimum element 0 and maximum element 1, and if we deliberately adjoin minimum and maximum elements $\pm\infty$ to the tropical semiring of real numbers then we obtain another complete idempotent semiring. Notice that each of these semirings has a negation-like operation which reverses the order of the elements. In the case of the Boolean semiring, negation is simply the involution interchanging 0 and 1, while in the case of the completed tropical semiring it is the involution sending each real number a to $-a$ and interchanging $\pm\infty$.

Cohen et al. [20] have established several Hahn-Banach type results about complete idempotent semimodules over a class of complete idempotent semirings that includes the Boolean semiring and the completed tropical semiring described above. Specifically, they demonstrate how to extend continuous linear functionals defined on complete subsemimodules of complete semimodules, as well as how to separate points from complete subsemimodules (see Lax [60] for an introduction to standard Hahn-Banach results in functional analysis). In the case of the completed tropical semiring, these results generalise the separation result of Hollings and Kambites [40] for row spaces of matrices. Similar functional analysis results in an idempotent setting have been obtained by Litvinov et al. [62], who also give a very thorough commentary on the history of such results.

The crucial common feature of the so-called reflexive semirings considered by Cohen et al. [20] is the presence of an order-reversing negation operation, which technically does not need to be an involution. In addition to their Hahn-Banach type extension and separation results, this assumption allowed them to construct a lattice anti-isomorphism between the row space and column space of each matrix with entries in a reflexive semiring. Hollings and Kambites [40] improved upon this result in the case of tropical semirings. Using the fact that tropical negation actually

is an involution, they showed that the anti-isomorphisms between row and column spaces of tropical matrices also have algebraic properties. Develin and Sturmfels [24] have shown that the same anti-isomorphisms induce combinatorial isomorphisms between the polyhedral complexes associated with the row and column spaces of each matrix.

As we mentioned above, an indispensable tool in tropical algebra is the notion of residuation. Loosely speaking, residuation is the ability to distinguish best, i.e., closest, approximate solutions to linear equations, and as such it gives us a way to carry out division (to a greater or lesser degree of accuracy). The study of residuated ordered algebraic structures was initiated in around 1920–40 by Krull [54] and Ward and Dilworth [84] in connection with lattices of ideals in commutative ring theory. Residuated structures, particularly residuated lattices, have since been adopted by researchers working in algebra, logic and data analysis, as they provide a setting in which to consider lattice-ordered groups, Boolean algebras, many-valued logics and fuzzy formal concepts. We will describe these applications shortly, but first we introduce residuated lattices.

A residuated lattice is essentially an idempotent semiring which is simultaneously a lattice, and in which each linear equation has a maximal approximate solution. In particular, this means that if a and b are elements of a residuated lattice then there is a maximal approximate solution to the linear equation $ax = b$, that is, there is a maximal element x satisfying $ax \leq b$. For example, the tropical semiring of real numbers, its completion with $\pm\infty$ adjoined, and the Boolean semiring are all residuated lattices. Galatos et al. [27] describe more residuated lattices; for a general introduction to residuation theory, see Blyth [9] or the foundational work of Blyth and Janowitz [10].

Initially developed by Wille [88] in the 1980s, formal concept analysis is a data analysis method that extracts certain concepts from the data describing the relationship between a given set of objects and a given set of attributes that the objects may or may not have. Specifically, a formal concept is a subset X of objects and a subset Y of attributes such that Y is the set of attributes that the objects X have in common and such that X is the set of objects exhibiting the attributes Y. The set of all formal concepts forms a lattice, so is accordingly called the concept lattice. A discussion of the foundations of formal concept analysis is given by Ganter and Wille [28].

In applications of formal concept analysis, the relationship between the objects and the attributes is often not binary; it may not be appropriate to say that a particular object has, or does not have, a particular attribute, but rather that there is only an extent to which the object has the attribute. Such concerns have led some researchers to reinterpret the theory of formal concepts in many-valued logics (see Bělohlávek [13]). The many-valued logics considered include intuitionistic logic and Łukasiewicz logic, which are modelled by the classes of Heyting algebras and MV-algebras respectively. We do not give the definitions of these algebras here (see Galatos et al. [27] for the relevant details), but the important point about them is that they are residuated lattices.

It turns out that the object-attribute data from which a formal concept is derived can be represented by a Boolean matrix with object many rows and attribute many columns, where each entry records whether a particular object has a particular attribute. More generally, if a many-valued logic is being used then this matrix is allowed to have entries in a (fixed, complete) residuated lattice, so that each entry records the extent to which a particular object has a particular attribute. Bělohlávek and Konečný [14] have studied the row and column spaces of matrices with entries in residuated lattices, and have generalised several results of Boolean matrix theory (see Kim [52]).

2 Basic definitions, notation and conventions

In this section we review the elementary concepts that will allow us to give a precise definition of what we mean by a 'semiring' (see Definition 4.1). We also remark on our notational and mathematical conventions. Finally, we recall standard notation for the semirings mentioned in section 1. An introduction to elementary order theory, including a discussion of lattices, will be given in section 15.

Definition 2.1 Let S be a set and let $\cdot$ be a binary operation on S. Then $(S, \cdot)$ is a *semigroup* if $a \cdot (b \cdot c) = (a \cdot b) \cdot c$ for all $a, b, c \in S$, that is, if $\cdot$ is *associative*.

If $(S, \cdot)$ is a semigroup then we will write ab instead of $a \cdot b$ for $a, b \in S$ unless it happens to be clearer to explicitly mention the binary operation. Furthermore, we will usually just say that S, rather than $(S, \cdot)$, is a semigroup if the binary operation is intended to be written in this way. If we have $ab = ba$ for all $a, b \in S$ then S is called *commutative*. If there is an *identity element* $1 \in S$ satisfying $a1 = a = 1a$

for all $a \in S$ then $(S, \cdot, 1)$ is called a *monoid*, and whenever $a, b \in S$ with $ab = 1$ we call a a *left inverse* of b and, dually, b a *right inverse* of a. A monoid in which each element a has a simultaneous left and right inverse a^{-1} is called a *group*.[1] An element $a \in S$ of a semigroup S is called *left cancellative* if $b = c$ whenever $b, c \in S$ with $ab = ac$. Dually, a is called *right cancellative* if $b = c$ whenever $b, c \in S$ with $ba = ca$, and if each element of S is both left and right cancellative then S is called *cancellative*. Notice that every group is cancellative. Standard introductions to the theory of semigroups include Clifford and Preston [17] and Howie [41].

Definition 2.2 Let S be a semigroup, let X be a set and let $\cdot\colon X \times S \to X$. Then $\cdot$ is a *right action* of S on X if $x \cdot (ab) = (x \cdot a) \cdot b$ for all $a, b \in S$ and all $x \in X$, that is, if $\cdot$ is *associative*.

As with the binary operation on a semigroup $(S, \cdot)$, we will use juxtaposition to indicate the action of S on a set X so that the condition in Definition 2.2 becomes $x(ab) = (xa)b$ for all $a, b \in S$ and all $x \in X$ (hence brackets are not necessary). Dually, a *left action* of S on X is defined to be an operation $\cdot\colon S \times X \to X$ satisfying $(ab)x = a(bx)$ for all $a, b \in S$ and all $x \in X$. Actions of semigroups on sets are also called 'S-acts', 'S-sets', 'S-operands', 'S-polygons', 'S-systems' and 'S-automata' in the literature (see Kilp et al. [51, page 43]). If S is a monoid then a right action of S on a set X satisfying $x1 = x$ for all $x \in X$ will be called a *right monoid action* and, dually, a left action of S satisfying $1x = x$ for all $x \in X$ will be called a *left monoid action*.

At the most fundamental level, nearly all of the results presented in this thesis concern actions of semigroups or monoids in one way or another, and this means that many results have a dual statement with the words 'left' and 'right' (or derived terms such as 'row space' and 'column space') interchanged. In the interest of brevity, we will usually neither prove nor even state dual versions of results, although we will make some dual definitions explicit at least. Our choice of which version of a result to give is informed by our right bias, which in turn stems from our notation for function application. If $f\colon X \to Y$ is a function between sets X and Y then fx denotes the element of Y obtained by applying f to x. This convention makes right actions the natural primary objects of study, as action-preserving functions then satisfy $f(xa) = (fx)a$ rather than the less mnemonic $f(ax) = a(fx)$ for left actions.

[1] It is actually enough for each element to have a left inverse (see Dickson [26]).

Throughout this thesis we write $\mathbf{N}$ for the natural numbers (starting with 1), $\mathbf{Z}$ for the integers, $\mathbf{R}$ for the real numbers and $\mathbf{C}$ for the complex numbers, each equipped with the usual operations of addition and multiplication. We also write $\mathbf{B}$ for the set $\{0, 1\}$ equipped with the operations of maximum and minimum, which we think of as "addition" and "multiplication" respectively. This structure is called the *Boolean semiring*. We mainly follow Hollings and Kambites [40] in our terminology and notation for the various tropical semirings, except we use 'max' and '+' instead of their special symbols '$\oplus$' and '$\otimes$'.

Definition 2.3

(i) The *finitary tropical semiring* $\mathbf{FT}$ is the set $\mathbf{R}$ equipped with the operations of maximum and addition.

(ii) The *tropical semiring* $\mathbf{T}$ is the set $\mathbf{R} \cup \{-\infty\}$ equipped with the operations of maximum and addition, where $-\infty \leq a$ and $a + (-\infty) = -\infty + a = -\infty$ for all $a \in \mathbf{R} \cup \{-\infty\}$.

(iii) The *completed tropical semiring* $\overline{\mathbf{T}}$ is the set $\mathbf{R} \cup \{-\infty, \infty\}$ equipped with the operations of maximum and addition, where $-\infty \leq a \leq \infty$ and $a + (-\infty) = -\infty + a = -\infty$ for all $a \in \mathbf{R} \cup \{-\infty, \infty\}$, and where $a + \infty = \infty + a = \infty$ for all $a \in \mathbf{R} \cup \{\infty\}$.

When first introducing a semiring, it is customary to distinguish its two identity elements. However, since our definition of a semiring does not require either identity element to exist, we need only specify the two binary operations that are to be taken as "addition" and "multiplication" (by convention, the first operation mentioned is intended to be addition). In section 4 we will present the formal definition of what we mean by a semiring, and we will justify why the semirings introduced above conform to that definition.

3 Summary of the thesis

In this thesis we study three related problems associated with linear algebra over arbitrary semirings, and we give solutions to these problems for certain rings and idempotent semirings. Many of our definitions and results build upon the work of Wilding et al. [86].

Some of the semirings that we are interested in do not possess an additive identity element (e.g., the finitary tropical semiring $\mathbf{FT}$) and some do not have a multiplicative identity element either (e.g., the semiring of 2×2 matrices with entries in $\mathbf{FT}$), so the first challenge is to formulate a workable theory of linear algebra over such semirings. The key observation is that it is sufficient for a semiring to have 'local' identity elements for each finite set of elements, rather than for the whole semiring. In section 4 we introduce precisely what we mean by a semiring with local identities, as well as the corresponding notion of a module over such a semiring, and we show that these definitions recover the standard ones when the whole semiring has an additive identity element and a multiplicative identity element. In section 5 we consider related concepts such as linear functions between modules, submodules, homomorphisms between semirings, subsemirings and semiring retracts. Again, we show that these concepts generalise their standard counterparts.

The presence of local identity elements allows us to define the row and column spaces of a matrix in an unambiguous way—either abstractly as modules of vectors, or concretely in terms of linear combinations and matrix multiplication. In section 6 we establish some basic properties of matrix multiplication and we show that Green's $\mathcal{L}$ and $\mathcal{R}$ relations for matrices can be characterised by row space and column space equality respectively (see Proposition 6.6). We also show that Green's $\mathcal{D}$ relation induces isomorphism of row/column spaces (see Proposition 6.7). In section 7 we discuss direct products and monoid semirings for finite monoids, and we show that if S is a semiring and G is a finite group then the group semiring SG is a retract of a full matrix semiring (see Theorem 7.3).

When presented with a semiring S, our first problem is to describe the 'kernels' of the row space and column space of each matrix with entries in S. The kernel of a set of vectors is an equivalence relation which generalises the orthogonal complement of a set, and which has been studied by Cohen et al. [18]. We consider kernels in section 8. Externally, the equivalence classes of the kernel of the row space of a matrix A form a module which is isomorphic to the column space of A (see Proposition 8.6), so the non-trivial aspect of the problem is to describe the internal structure of each kernel class. As an aid to doing this we introduce the kernel of a relation, and it turns out that together these two notions of kernel constitute a Galois connection between sets of vectors and relations on vectors (see Proposition 8.4). Cohen et al. [19] have also considered kernels of relations in the case of the completed tropical semiring $\overline{\mathbf{T}}$

as a way to talk about 'separability'. Accordingly, we define separability in terms of the closed elements of the Galois connection just mentioned.

Our second problem is to determine which linear functionals on the row space and column space of each matrix over a semiring S have a linear extension defined on the appropriate containing module of vectors. The ideal situation is that every such function has an extension (a restricted form of self-injectivity of S), in which case S is called 'exact'. In section 9 we characterise extendibility of linear functionals in terms of representability by vectors (see Proposition 9.6) and, subsequently, we show that exactness can be rephrased as a property of the kernel Galois connection discussed above (see Proposition 9.10). We also show that the row and column spaces of any matrix over an exact semiring are isomorphic to each others' duals, and that exactness allows us to characterise Green's $\mathcal{D}$ relation for matrices by isomorphism of row/column spaces (see Theorem 9.11).

As we discussed in section 1, the row and column spaces of a matrix with entries in a field are isomorphic because they are vector spaces of the same dimension, but isomorphism is not necessarily the "natural" choice of relationship between the row and column spaces of a matrix over an arbitrary semiring. One reason to believe this is that the row space of a matrix is naturally a left module, while the column space is naturally a right module, and so there ought not to be an isomorphism between the two modules. More plausibly, the row and column spaces of each matrix over $\mathbf{FT}$ are naturally "anti-isomorphic" in the sense of section 16, not isomorphic, and for matrices over $\mathbf{C}$ both isomorphism and conjugate isomorphism are equally natural. Our third problem is therefore to explain the ways in which the row and column spaces of each matrix over a semiring can be related; when presented with a semiring we would ideally like to know which forms of isomorphism (if any) to naturally expect between row and column spaces. In section 10 we introduce the notion of a 'conjugation' on a semiring in an attempt to address this problem, and we show that complex conjugation on $\mathbf{C}$ is such a conjugation (see Theorem 10.4).

We initiate detailed study of exactness in section 11 by showing that if S is an exact semiring then each full matrix semiring over S is exact (see Proposition 11.1), and that the product of exact semirings is exact (see Proposition 11.2). In section 12 we then consider exactness of subsemirings and subsemirings of exact semirings. We show that if S is an exact subsemiring of a semiring T then the finitely generated ideal structure of T must be at least as complicated as the finitely generated ideal

structure of S (see Theorem 12.2). Going in the other direction, we give sufficient conditions for a retract of an exact semiring to be exact (see Theorem 12.4), and we apply this result to show that if S is an exact semiring and G is a finite group then the group semiring SG is exact (see Corollary 12.5). We also show that if S has identity elements and is exact then essentially the only 'shapes' of matrix semiring over S that are exact are the symmetric ones, e.g., semirings of diagonal matrices, not semirings of upper triangular matrices (see Theorem 12.3).

In passing, we pose a Baer type question in section 12: is a semiring S exact if all linear functionals on finitely generated ideals of S have linear extensions defined on S? This question remains open when restricted to rings (it is unknown whether every 'F-injective' ring is 'FP-injective'; see Nicholson and Yousif [70, Question 10]), but it has a negative answer as stated for semirings. Indeed, in section 20 we construct a non-exact idempotent semiring with the property that each linear functional on a finitely generated ideal has an extension (see Example 20.6).

In section 13 we show that exactness for rings (also known as FP-injectivity) is characterised by a double orthogonal complement condition on the row and column spaces of matrices (see Proposition 13.3). Several classes of rings are already known to be exact, so it would not be especially productive to attempt to verify this double orthogonal complement condition directly. Instead, we introduce the stronger notion of an 'exact annihilator ring', over which the row space of each matrix is the orthogonal complement of some column space (and vice versa). Exact annihilator rings are desirable because the matrices used to "witness" exactness also provide a description of orthogonal complements (see Proposition 13.8). In section 14 we demonstrate how to construct such witnesses for matrices over Boolean rings and proper homomorphic images of principal ideal domains. The fact that every Boolean ring is an exact annihilator ring means that there exist exact rings which are not self-injective (see page 89). This answers a question of Wilding et al. [86].

As well as being exact annihilator rings, Boolean rings and proper homomorphic images of principal ideal domains are commutative elementary divisor rings. The orthogonal complements of the row and column spaces of a matrix over such a ring can be described as quotients (see Theorem 14.3), so in this regard Boolean rings and proper homomorphic images of principal ideal domains behave just as if they were fields.[1] We also show in section 14 that the identity function on any commutative

[1] Some of them are in fact fields.

elementary divisor ring (e.g., $\mathbf{Z}$) is a conjugation in the sense outlined above, and consequently the row space and column space of each matrix over such a ring are isomorphic modules (see Theorem 14.2).

The other type of semiring we are interested in is residuated lattices. These semirings are best studied from within algebraic order theory, so in sections 15 and 16 we provide an introduction to the prerequisite theory of partial orders and ordered algebraic structures respectively. The discussion in section 15 covers posets, lattices, Boolean algebras, monotone functions, adjunctions, order isomorphisms, antitone functions, Galois connections and order anti-isomorphisms. The discussion in section 16 covers ordered monoids, their actions on posets and the appropriate notions of structure preserving functions and isomorphisms. We also introduce an equivalent formulation of isomorphisms, which leads us to a sensible and powerful definition of what we mean by an anti-isomorphism in this setting.

In section 17 we recall the definitions of residuated monoids and residuated lattices, along with some standard illustrative examples. We also introduce residuation in the more general context of ordered monoid actions. The residuated structures we consider here can be interpreted as categories enriched over the acting monoid M (see Proposition 17.3), so it is natural to ask for a characterisation of the categories enriched over M that arise in this way. As a partial answer, we give conditions for an enriched category to have a residuated structure as a quotient (see Theorem 17.4).

In section 18 we consider linear algebra over residuated lattices, which can be treated as semirings by taking addition to be the lattice join operation. After establishing that residuated lattices have local identity elements (see Proposition 18.1), we describe how to extend residuation to matrices via the lattice meet operation (see Proposition 18.2). In particular cases this observation is not new; Cohen et al. [19] have described the same residuation operations for matrices over $\overline{\mathbf{T}}$, while Bělohlávek and Konečný [14] have defined similar operations for matrices with entries in complete commutative residuated lattices. Matrix residuation can be used to show that row and column spaces are always lattices (see Proposition 18.3). It also allows us to understand the internal structure of kernel classes (see Proposition 18.5).

A residuated lattice viewed as a semiring may or may not be exact, so if we want to obtain exactness then we must impose some additional structure. In section 19 we restrict to the case where residuation emerges from the presence of an (order reversing) involution, and we show that all such residuated lattices are exact (see

Theorem 19.4). Our approach essentially follows Cohen et al. [20, Theorem 34], except we permit non-complete residuated lattices (e.g., **FT**). We also show that the involution is a conjugation which induces an anti-isomorphism between the row space and column space of each matrix (see Theorem 19.5). The anti-isomorphisms we obtain here are more algebraic in flavour than the ones obtained by Cohen et al. [20, Theorem 42] because we assume that the underlying residuated lattice has an involution, whereas Cohen et al. [20] make a slightly weaker assumption. The results in section 19 apply to the semirings **B**, **FT** and $\overline{\mathbf{T}}$, but not to **T**.

In section 20 we study powersets of monoids, viewed as residuated lattices (and hence as semirings). We begin by showing that the powerset of a monoid is involutive in the above sense if and only if the monoid is a group (see Theorem 20.1). This result implies that the semiring of subsets of an arbitrary group is exact, but it does not rule out the possibility that the semiring of subsets of a non-group monoid might be exact for some other reason. If such a non-group monoid exists then it cannot be cancellative (see Proposition 20.3). We also show that the powerset of a finite monoid is exact if and only if the monoid is a group (see Corollary 20.5), so if there does exist a non-cancellative monoid whose powerset is exact then it must be infinite. The existence of such a monoid remains an open question deserving of further investigation.

Semirings, modules and matrices

4 Semirings and modules

As we discussed in section 3, we would like to be able to study linear algebra over a variety of semirings, including certain semirings that are missing one or both of the usual identity elements. With care, we can (for the most part) treat such semirings just as we would any other semiring once we have established that linear algebra still "works" in the absence of identity elements. This does require some non-standard definitions, however, along with proofs of facts that are usually automatic, so in this and the next section we prepare the foundational definitions and results that are needed to obtain basic, expected, properties of semirings, modules and matrices.

The fundamental observation, made by Hollings and Kambites [40], is that in most circumstances it is good enough to have elements that behave like identity elements "locally", but that are not necessarily true (i.e., "global") identity elements. In the context of linear algebra, local means finite because matrices have only finitely many entries. Put another way, an application of global identity elements—such as multiplication by an identity matrix—can often be replaced by an application of identity elements local to the entries of some matrix. This observation motivates the following precise definition of what we mean by a semiring with local identities.

Definition 4.1 Let S be a non-empty set, let $+$ be a binary operation on S with $(S, +)$ a commutative semigroup and let $\cdot$ be a binary operation on S with $(S, \cdot)$ a semigroup. Then $(S, +, \cdot)$ is a *semiring* if

(i) $c(a + b) = ca + cb$ and $(a + b)c = ac + bc$ for all $a, b, c \in S$; and if

(ii) for each non-empty finite $L \subseteq S$ there are *local identities* $0_L, 1_L \in S$ satisfying $a + b0_L = a1_L = a$ and $a + 0_L b = 1_L a = a$ for all $a, b \in L$.

As usual, we will simply say that S, rather than $(S, +, \cdot)$, is a semiring unless we want to draw attention to the operations that are to be used as addition and

multiplication. Notice that addition on a semiring S is defined to be commutative, whereas multiplication need not be. We will say that S is a *commutative* semiring in the case multiplication on S is commutative, although there will be very few instances where we will actually need this additional assumption.

Hollings and Kambites [40] say that a (commutative) semiring S has 'local zeros' if for each non-empty finite $L \subseteq S$ there is some $c \in S$ with $a + c = a$ for all $a \in L$. Definition 4.1 (ii) clearly implies this condition, as we can just take $c = b0_L$ for any $b \in L$, but our condition captures more of what it means for an element to behave like a global zero would. Specifically, condition (ii) tries to mimic the absorbing nature of zero by insisting that when 0_L is multiplied by any $b \in S$, the resulting elements $b0_L$ and $0_L b$ are still zero-like enough in the additive sense. We also impose a less complicated "local ones" condition that requires each non-empty finite $L \subseteq S$ to have an associated element 1_L that, on L, behaves exactly like a multiplicative identity element would. As the following example demonstrates, Hollings and Kambites [40] have no need for such a condition because the semirings they consider already have multiplicative identity elements.

Example 4.2 The finitary tropical semiring $\mathbf{FT} = (\mathbf{R}, \max, +)$ is a commutative semiring in the sense of Definition 4.1. If $L \subseteq \mathbf{FT}$ is non-empty and finite then

$$0_L = \min\{a - b : a, b \in L\} \tag{4.1}$$

and

$$1_L = 0 \tag{4.2}$$

satisfy $\max\{a, b + 0_L\} = a + 1_L = a$ for all $a, b \in L$ (see page 108 for 0_L), and as such Definition 4.1 (ii) holds for $\mathbf{FT}$. In other words, $\mathbf{FT}$ has local identities.

The restriction to non-empty subsets in Definition 4.1 (ii) is not strictly necessary, as if S is a semiring we can certainly find "local identities" $0_\emptyset, 1_\emptyset \in S$ satisfying the vacuous conditions $a + b0_\emptyset = a1_\emptyset = a$ and $a + 0_\emptyset b = 1_\emptyset a = a$ for all $a, b \in \emptyset$; simply take $0_\emptyset$ and $1_\emptyset$ to be any elements of S. However, since such elements $0_\emptyset$ and $1_\emptyset$ are devoid of meaning, we prefer to overlook their existence. Moreover, removing the supposition of non-emptiness in Definition 4.1 (ii) would complicate matters when defining local identities in a uniform way. For example, in Proposition 18.1 we define local identities for each finite subset L of a residuated lattice in a way that (by convention) specifies a possibly non-existent top element in the case $L = \emptyset$.

A semiring which has an additive identity element will be called a 0-*semiring*, and similarly a semiring which has a multiplicative identity element will be called a 1-*semiring*. For example, **FT** is a 1-semiring with multiplicative identity element $0 \in \mathbf{R}$, but is not a 0-semiring because it has no bottom element.[1] That is, there is no element $\perp \in \mathbf{R}$ satisfying $\max\{a, \perp\} = a$ for all $a \in \mathbf{R}$. The tropical semiring $\mathbf{T} = (\mathbf{R} \cup \{-\infty\}, \max, +)$ is a 0-semiring though, because the adjoined bottom element $-\infty$ satisfies $\max\{a, -\infty\} = a$ for all $a \in \mathbf{T}$.

A semiring which is both a 0-semiring and a 1-semiring will be called a *standard semiring*. Despite the name, it is not yet clear that this definition actually matches what Golan [34] calls a semiring, for although a standard semiring S automatically comprises an additive commutative monoid $(S, +, 0)$ and a multiplicative monoid $(S, \cdot, 1)$, Golan [34, page 1] also requires $0 \in S$ to satisfy $a0 = 0 = 0a$ for all $a \in S$. However, the following result shows that this property follows from Definition 4.1 in the case S is a 0-semiring, and thus a standard semiring really is deserving of that title.

Proposition 4.3 *If S is a 0-semiring then $a0 = 0 = 0a$ for all $a \in S$.*

Proof Let $a \in S$ and take $L = \{0, a0\} \subseteq S$. Then by Definition 4.1 (ii) there is some $0_L \in S$ satisfying $0 + 00_L = 0$ and $0 + a00_L = 0$, and thus we have $00_L = 0$ and $a00_L = 0$ because 0 is the additive identity element in S. Therefore $a0 = a00_L = 0$, as required; a dual argument confirms that $0a = 0$. □

As well as **T** and the completed tropical semiring $\overline{\mathbf{T}} = (\mathbf{R} \cup \{-\infty, \infty\}, \max, +)$, other examples of standard semirings include **Z**, **R** and **C** because they each contain the usual identity elements. The Boolean semiring $\mathbf{B} = (\{0, 1\}, \max, \min)$ is also a standard semiring, with additive identity element 0 and multiplicative identity element 1. However, as we have defined it in section 2, **N** is not a standard semiring because it does not contain $0 \in \mathbf{Z}$. In fact, **N** is not strictly a semiring in the sense of Definition 4.1, as it does not contain a local identity $0_{\{1\}}$ either.

If each element a of a standard semiring R has an additive inverse $-a \in R$, so that the commutative monoid $(R, +, 0)$ is a group, then R is called a *ring* (see Cohn [21, page 21]). In this case a standard argument using the absorbing property of $0 \in R$ shows that we have $a(-1) = -a = (-1)a$ for all $a \in R$. Of the examples mentioned above, only **Z**, **R** and **C** are rings. One of the key tools in the study of rings is

[1] Yes, this is confusing!

the notion of a 'module' over a ring—essentially a commutative group on which the ring acts—and, because of its usefulness, this idea has also been considered for semirings. To account for local identities, our definition of a module over a semiring differs slightly from the standard definition discussed below.

Definition 4.4 Let S be a semiring, let $(X, +)$ be a commutative semigroup and let $\cdot$ be a right action of $(S, \cdot)$ on X. Then $(X, +, \cdot)$ is a *right S-module* if

(i) $x(a + b) = xa + xb$ and $(x + y)a = xa + ya$ for all $a, b \in S$ and all $x, y \in X$; and if

(ii) for each non-empty finite $L \subseteq X$ there are *right local identities* $0_L, 1_L \in S$ satisfying $x + y0_L = x1_L = x$ for all $x, y \in L$.

The similarity between Definitions 4.1 and 4.4 makes it obvious that if S is a semiring then S is itself a right S-module; the right action of S on S is just given by multiplication. The notion of a *left S-module* is defined dually, and it is similarly clear that S is always a left S-module. Proposition 5.6 generalises these observations.

Golan [34, page 149] defines a 'semimodule' over a standard semiring S to be a commutative monoid $(X, +, 0)$ on which $(S, \cdot)$ acts, such that Definition 4.4 (i) holds, and such that

(i) $x0 = 0 = 0a$ for all $a \in S$ and all $x \in X$; and

(ii) $x1 = x$ for all $x \in X$.

We do not impose these last two conditions because we only require the existence of local identities, but, just as with semirings, Definition 4.4 (ii) is formulated so as to ensure that if S does happen to be a standard semiring then these two conditions automatically hold. Proposition 4.5, below, confirms this. Moreover, if R is a ring then then our R-modules are just modules in the conventional sense (see Cohn [21, page 226]) because each $x \in X$ has an additive inverse $-x \in X$ defined by $-x = x(-1)$. It is for this reason that we use the term 'module'. Note, however, that Golan [34, page 150] reserves this term for a semimodule in which each element has an additive inverse—even if the underlying standard semiring is not a ring.

Proposition 4.5 *Let S be a semiring and let X be a right S-module.*

(i) *If S is a 0-semiring then X has an additive identity element $0 \in X$ satisfying $x0 = 0 = 0a$ for all $a \in S$ and all $x \in X$.*

(ii) *If S is a 1-semiring then $x1 = x$ for all $x \in X$.*

Proof (i). We will first show that $y0$ is the additive identity element in X for any chosen $y \in X$. Let $x \in X$ and take $L = \{x, y0\}$. Then by Definition 4.4 (ii) there is some $0_L \in S$ satisfying $x + y00_L = x$. By Proposition 4.3 we have $0 = 00_L$, so $x + y0 = x$, and as such $y0$ is the additive identity element in X. Identity elements are unique, so $y0$ does not depend on y, and this means that $x0 = y0$ for all $x \in X$. It therefore remains to show that $y0 = y0a$ for all $a \in S$, but this holds because we have $0 = 0a$ by Proposition 4.3 again.

(ii). Let $x \in X$ and take $L = \{x\}$. Then by Definition 4.4 (ii) there is some $1_L \in S$ satisfying $x1_L = x$. Hence $x1 = x1_L1 = x1_L = x$ because $1_L1 = 1_L$. □

5 Linear functions and homomorphisms

Given a fixed semiring S, structure preserving functions between S-modules ought to respect the operations that Definition 4.4 requires an S-module to have. Specifically, we require a structure preserving function between right S-modules to respect addition and the two right actions of S in the following sense.

Definition 5.1 Let S be a semiring and let $f\colon X \to Y$ be a function between right S-modules X and Y. Then f is *right S-linear* if

(i) $f(x + y) = fx + fy$ for all $x, y \in X$; and if

(ii) $f(xa) = (fx)a$ for all $a \in S$ and all $x \in X$.

Our definition of a linear function between S-modules is identical to the standard definition of a 'homomorphism' between S-modules (see Golan [34, page 156]). This might come as a surprise because Definitions 4.1 and 4.4 differ from the standard definitions of a semiring and a module, respectively, so we might have expected Definition 5.1 to also impose a local identities condition on a linear function. The reason no such condition is required is that local identities for an S-module lie in the semiring S, not the S-module, and thus a linear function between S-modules does not directly interact with local identities for the S-modules.

If X and Y are right S-modules with $X \subseteq Y$ and the inclusion function $X \hookrightarrow Y$ right S-linear then we will say that X is a *right S-submodule* of Y. In such a case Definition 5.1 (i) ensures that addition on X is just the restriction of addition on

Y, and similarly Definition 5.1 (ii) ensures that the action of S on X is just the restriction of the action of S on Y. That is, X is closed under addition on Y and the action of S on Y, and in fact this is all we need to check in order for a non-empty subset of Y to be an S-submodule of Y. In particular, when trying to show that a non-empty subset X of a right S-module Y is actually a right S-submodule of Y, we do not need to worry about verifying Definition 4.4 (ii) because local identities for $L \subseteq X$ will be the same as for $L \subseteq Y$ provided X is closed under addition and the right action of S. This mirrors what happens for standard semirings and modules (see Golan [34, page 150]).

The image of a linear function is an important instance of a submodule. If $f\colon X \to Y$ is a right S-linear function between right S-modules X and Y then $\mathrm{Im}(f) = \{fx : x \in X\}$ is a right S-submodule of Y because, by Definition 5.1, it is closed under addition and the right action of S.

The other important object associated with f is its set-theoretic kernel, i.e., the equivalence relation defined by $\mathrm{Ker}(f) = \{(x, x') \in X \times X : fx = fx'\}$. In general we cannot define the kernel of f to be the set $\{x \in X : fx = 0\}$ because Y does not necessarily contain an additive identity element 0. Moreover, even if Y does have a 0 element, this alternative definition of $\mathrm{Ker}(f)$ would not be guaranteed to provide useful information about f; for $\{x \in X : fx = 0\}$ to always give as much information as $\mathrm{Ker}(f)$ we would need X and Y to have additive inverses as well.

Although $\mathrm{Ker}(f)$ it is not a submodule of X, it does at least have some interaction with the right S-module structure of X. Specifically, $\mathrm{Ker}(f)$ is a *right S-congruence* on X. This means that $\mathrm{Ker}(f)$ is compatible with addition in the sense that if $(x, x'), (y, y') \in \mathrm{Ker}(f)$ then $(x + y, x' + y') \in \mathrm{Ker}(f)$, and that it is compatible with the right action of S in the sense that if $a \in S$ and $(x, x') \in \mathrm{Ker}(f)$ then $(xa, x'a) \in \mathrm{Ker}(f)$. These properties make it possible to turn the set $X/\mathrm{Ker}(f)$ of equivalence classes of $\mathrm{Ker}(f)$ into a right S-module by defining

$$[x]_{\mathrm{Ker}(f)} + [y]_{\mathrm{Ker}(f)} = [x + y]_{\mathrm{Ker}(f)} \tag{5.1}$$

and

$$[x]_{\mathrm{Ker}(f)}a = [xa]_{\mathrm{Ker}(f)} \tag{5.2}$$

for all $a \in S$ and all $x, y \in X$. Note that if $L \subseteq X/\mathrm{Ker}(f)$ is non-empty and finite then to satisfy Definition 4.4 (ii) we can simply take $0_L = 0_K$ and $1_L = 1_K$ where

K is any set of representatives of the classes in L. That is, right local identities for $X/\operatorname{Ker}(f)$ are inherited from Y.

The kernel and image of a right S-linear function are connected by the following "first isomorphism theorem", where, as usual, an *isomorphism* of right S-modules is defined to be a right S-linear bijection. By a standard argument, the inverse of an isomorphism is also right S-linear, so we may say that two right S-modules are isomorphic (written $X \cong Y$) without specifying a direction.

Proposition 5.2 *If $f\colon X \to Y$ is a right S-linear function between right S-modules X and Y then $X/\operatorname{Ker}(f) \cong \operatorname{Im}(f)$ as right S-modules.*

Proof It follows immediately from the above definition of $\operatorname{Ker}(f)$ that the surjective function $\operatorname{Im}(f) \to X/\operatorname{Ker}(f)$ given by $fx \mapsto [x]_{\operatorname{Ker}(f)}$ is well-defined and injective. Moreover, together with (5.1) and (5.2), right S-linearity of f ensures that this bijection $\operatorname{Im}(f) \to X/\operatorname{Ker}(f)$ is also right S-linear, so is an isomorphism of right S-modules. □

We will mainly be interested in studying functions between modules, but it will occasionally be useful to consider functions between semirings. Unlike linear functions, however, our definition of a structure preserving function between semirings is quite complicated because it must account for local identities.

Definition 5.3 Let $f\colon S \to T$ be a function between semirings S and T. Then f is a *homomorphism* if

(i) $f(a+b) = fa + fb$ for all $a, b \in S$;

(ii) $f(ab) = (fa)(fb)$ for all $a, b \in S$; and if

(iii) for each non-empty finite $L \subseteq T$ there are *local identities* $0_L, 1_L \in \operatorname{Im}(f)$ satisfying $a + b0_L = a1_L = a$ and $a + 0_L b = 1_L a = a$ for all $a, b \in L$.

Conditions (i) and (ii) are unsurprising; they merely say that a homomorphism must respect the semiring operations. The motivation for (iii) is much less obvious, even though it clearly has something to do with local identities. We now explain where this condition comes from, and why it is appropriate.

If S and T are standard semirings then a homomorphism $f\colon S \to T$ ought to satisfy $f0 = 0$ and $f1 = 1$ (see Golan [34, page 105]), but in the context of local

identities this condition splits into a "forwards" requirement and a "backwards" requirement. Firstly, local identities for S should remain local identities in the image of f, and secondly, there should be local identities for T that lie in the image of f. The first (forwards) requirement is taken care of by Definition 5.3 (i) and (ii): if L is a non-empty finite subset of S then we have

$$fa + (fb)(f0_L) = (fa)(f1_L) = fa \tag{5.3}$$

and

$$fa + (f0_L)(fb) = (f1_L)(fa) = fa \tag{5.4}$$

for all $a, b \in L$, and as such $f0_L$ and $f1_L$ still behave like local identities. The second (backwards) requirement is precisely what Definition 5.3 (iii) demands: our definition of a semiring requires T to have local identities, but Definition 5.3 (iii) goes one step further and insists that local identities can actually be taken to lie in $\mathrm{Im}(f)$.[1]

It is worth remarking that if f is surjective then Definition 5.3 (iii) is identical to Definition 4.1 (ii) for T, and so a surjective function between semirings is a homomorphism if and only if it respects addition and multiplication. This matches up with the standard situation, as it can be shown that a surjective function f between standard semirings satisfies $f0 = 0$ and $f1 = 1$ if it respects addition and multiplication. The following result gives further confirmation that our definition of a homomorphism is sensible.

Proposition 5.4 *Let $f\colon S \to T$ be a homomorphism between semirings S and T.*

(i) *If S and T are 0-semirings then $f0 = 0$.*

(ii) *If S and T are 1-semirings then $f1 = 1$.*

Proof (i). Take $L = \{0, f0\} \subseteq T$. Then by Definition 5.3 (iii) there is some $0_L \in \mathrm{Im}(f)$ satisfying $0 + (f0)0_L = 0$. That is, $(f0)0_L = 0$. Since $0_L \in \mathrm{Im}(f)$ there is some $a \in S$ with $0_L = fa$, and thus Proposition 4.3 and Definition 5.3 (ii) give $f0 = f(0a) = (f0)(fa) = 0$.

[1] Definition 5.3 (iii) does not say that all local identities for T must lie in $\mathrm{Im}(f)$, just that for each suitable L we can find some that do. These local identities 0_L and 1_L need not be the same ones first used to verify that T is a semiring, but sensibly chosen local identities will usually turn out to be in $\mathrm{Im}(f)$ anyway.

(ii). Take $L = \{1\} \subseteq T$. Then by Definition 5.3 (iii) there is some $1_L \in \mathrm{Im}(f)$ satisfying $1(1_L) = 1$. Therefore $1_L = 1$, and so since $1_L \in \mathrm{Im}(f)$ there is some $a \in S$ with $1 = fa$. Hence

$$f1 = (f1)1 = (f1)(fa) = f(1a) = fa = 1 \tag{5.5}$$

by Definition 5.3 (ii). □

We will say that a semiring S is a *subsemiring* of a semiring T if $S \subseteq T$ and the inclusion function $S \hookrightarrow T$ is a homomorphism. Proposition 5.4 ensures that if S and T are standard semirings with S a subsemiring of T then S and T share the same identity elements, so in this case our definition of a subsemiring is equivalent to the standard definition (see Golan [34, page 3]). What is not clear, however, is whether our subsemiring relation is transitive in general. That is, if S is a subsemiring of T and T is a subsemiring of U, is S a subsemiring of U? This question arises because we have not yet shown—and it is certainly not obvious—that the composition of two homomorphisms is again a homomorphism. The following result redresses this.

Proposition 5.5 *If $f\colon S \to T$ and $g\colon T \to U$ are homomorphisms between semirings S, T and U then $g \circ f$ is a homomorphism.*

Proof It is obvious that $g \circ f$ respects addition and multiplication, so it remains to show that for each non-empty finite $L \subseteq U$ there are $0_L, 1_L \in \mathrm{Im}(g \circ f)$ satisfying $a + b0_L = a1_L = a$ and $a + 0_L b = 1_L a$ for all $a, b \in L$. To simplify matters, we will only show that there are 0_L and 1_L satisfying $a + b0_L = a1_L = a$; the proof that they also satisfy the other requirement is dual.

Let $L \subseteq U$ be non-empty and finite. Then by Definition 5.3 (iii) for g there are $0_g, 1_g \in \mathrm{Im}(g)$ satisfying $a + b0_g = a1_g = a$ for all $a, b \in L$. Since $0_g, 1_g \in \mathrm{Im}(g)$ we can write $0_g = g0_T$ and $1_g = g1_T$ for some $0_T, 1_T \in T$. Now since $\{0_T, 1_T\}$ is a non-empty finite subset of T, Definition 5.3 (iii) for f tells us that there are $0_f, 1_f \in \mathrm{Im}(f)$ with $0_T + 1_T 0_f = 0_T$ and $1_T 1_f = 1_T$. Finally, take $0_L = g0_f$ and $1_L = g1_f$, which lie in the image of $g \circ f$ because $0_f, 1_f \in \mathrm{Im}(f)$.

To show that $a + b0_L = a$ for all $a, b \in L$, first observe that

$$0_g + 1_g 0_L = g0_T + (g1_T)(g0_f) = g(0_T + 1_T 0_f) \tag{5.6}$$

by Definition 5.3 (i) and (ii) for g. The fact that $0_T + 1_T 0_f = 0_T$ then gives

$$b(0_g + 1_g 0_L) = b(g0_T) = b0_g \tag{5.7}$$

for all $b \in L$, and thus if $b \in L$ then $b0_g + b0_L = b0_g$ because $b1_g = b$. Therefore

$$a + b0_L = a + b0_g + b0_L = a + b0_g = a \tag{5.8}$$

for all $a, b \in L$, as required.

To show that $a1_L = a$ for all $a \in L$, first observe that

$$1_g 1_L = (g1_T)(g1_f) = g(1_T 1_f) \tag{5.9}$$

by Definition 5.3 (ii) for g. The fact that $1_T 1_f = 1_T$ then gives

$$a1_g 1_L = a(g1_T) = a1_g \tag{5.10}$$

for all $a \in L$. Hence if $a \in L$ then $a1_L = a$, as required, because $a1_g = a$. □

Proposition 5.6 *Let $f\colon S \to T$ be a homomorphism between semirings S and T, and let X be a right T-module. Then X is a right S-module via the action of S given by $xa = x(fa)$ for all $a \in S$ and all $x \in X$.*

Proof The corresponding result for rings is well-known and is sometimes called 'restriction of scalars' (see Sharp [75, Remark 6.6]). In the context of semirings, the only non-routine part of proving that X is a right S-module is establishing the existence of right local identities that lie in S.

Let $L \subseteq X$ be non-empty and finite. Then by Definition 4.4 (ii) for X as a right T-module there $0_T, 1_T \in T$ satisfying $x + y0_T = x1_T = x$ for all $x, y \in L$. Now (as in the proof of Proposition 5.5) since $\{0_T, 1_T\}$ is a non-empty finite subset of T, Definition 5.3 (iii) tells us that there are $0_f, 1_f \in \mathrm{Im}(f)$ with $0_T + 1_T 0_f = 0_T$ and $1_T 1_f = 1_T$. The fact that $0_f, 1_f \in \mathrm{Im}(f)$ means that we can then write $0_f = f0_L$ and $1_f = f1_L$ for some $0_L, 1_L \in S$.

If $x, y \in L$ then

$$y0_T + y0_f = y0_T + y1_T 0_f = y(0_T + 1_T 0_f) = y0_T \tag{5.11}$$

because $y1_T = y$ and $0_T + 1_T 0_f = 0_T$, and thus

$$x + y(f0_L) = x + y0_f = x + y0_T + y0_f = x + y0_T = x \tag{5.12}$$

because $x + y0_T = x$. In terms of the right action of S on X defined above, this means that $x + y0_L = x$ for all $x, y \in L$. To show that $x1_L = x$ as well, let $x \in L$ and observe that $x1_T 1_f = x1_T$ because $1_T 1_f = 1_T$. The fact that $x1_T = x$ then gives

$$x(f1_L) = x1_f = x1_T 1_f = x1_T = x, \tag{5.13}$$

and as such $x1_L = x$. Hence Definition 4.4 (ii) is satisfied for X with the right action of S defined above. □

Proposition 5.6 is especially useful for subsemirings, as it tells us that if S is a subsemiring of T then every T-module can be viewed as an S-module. In particular, T itself can be viewed as an S-module for each subsemiring $S \subseteq T$, and this lets us consider S-linear functions between S and T. An important instance for us will be when there is a surjective S-linear function $T \to S$ that "collapses" T onto S, allowing us to transfer various properties of T to S (see Theorems 7.3 and 12.4). To make this condition precise, we introduce the following notion of a 'retract'. Our use of this term is consistent with its uses in topology (see Hatcher [36, page 3]) and group theory (see Lyndon and Schupp [63, page 2]).

Definition 5.7 Let S and T be semirings. Then S is a *right retract* of T if S is subsemiring of T and there is a right S-linear function $f\colon T \to S$ that fixes S pointwise, i.e., with $fa = a$ for all $a \in S$.

6 Matrices and Green's relations

Although we have introduced a completely abstract definition of a module over a semiring (see Definition 4.4), we are mainly interested in finitely generated modules of row vectors and column vectors. This is because our ultimate aim is to understand (or at least provide general strategies for understanding) the behaviour of matrices over a given semiring S, and two very important modules associated with a matrix are the left S-module generated by its rows (its 'row space') and the right S-module generated by its columns (its 'column space'). In this section we define all these

concepts, and we establish some basic properties of matrices that are usually taken for granted over standard semirings. We also link Green's relations for matrices with row spaces and column spaces.

Definition 6.1 Let S be a semiring and let $m, n \in \mathbf{N}$. A *matrix* of size $m \times n$ with entries in S is a rectangular array A of elements of S, arranged into m rows and n columns. More formally, A is simply a function $\{1, \ldots, m\} \times \{1, \ldots, n\} \to S$.

Given a semiring S, we write $S^{m \times n}$ for the set of $m \times n$ matrices with entries in S, and if $A \in S^{m \times n}$ is such a matrix then we write A_{ij} for the i-jth entry of S. That is, $A_{ij} \in S$ denotes the element at the ith position in column j (or equally the jth position in row i) of A, where $1 \leq i \leq m$ and $1 \leq j \leq n$. Special cases of Definition 6.1 occur when $m = 1$ and when $n = 1$. In the first case we call an element of $S^{1 \times n}$ a *row vector*, and, dually, in the second case we call an element of $S^{m \times 1}$ a *column vector*. Notice that in the case $m = n = 1$ the "matrices" in $S^{1 \times 1}$ are essentially the elements of S, so we just identify $S^{1 \times 1}$ with S.

Matrices of the same size can be summed entrywise, making each $\big(S^{m \times n}, +\big)$ a commutative semigroup. Specifically, if $A, B \in S^{m \times n}$ then we define another matrix $A + B \in S^{m \times n}$ by

$$(A + B)_{ij} = A_{ij} + B_{ij} \tag{6.1}$$

for all $1 \leq i \leq m$ and all $1 \leq j \leq n$. Matrices of appropriate sizes can also be multiplied: if $A \in S^{m \times n}$ and $B \in S^{n \times q}$ then we define a matrix $AB \in S^{m \times q}$ by

$$(AB)_{ik} = \sum_{j=1}^{n} A_{ij} B_{jk} \tag{6.2}$$

for all $1 \leq i \leq m$ and all $1 \leq k \leq q$. It is straightforward to check that this operation is associative whenever it is defined, so in particular each $\big(S^{n \times n}, \cdot\big)$ is a semigroup. Finally, matrix multiplication distributes over matrix addition in the sense of Definition 4.1 (i), and thus we are nearly able to say that each $\big(S^{n \times n}, +, \cdot\big)$ is a semiring.

If S is a standard semiring then the constructions and properties described above make each $S^{n \times n}$ a standard semiring (see Golan [34, page 27]), but if we are to conclude that $S^{n \times n}$ is a semiring even when S is not standard then we need to verify Definition 4.1 (ii) for $S^{n \times n}$. The following result does this in a way that is useful more generally.

Proposition 6.2 *Let S be a semiring and let $L \subseteq S^{m\times n}$ be non-empty and finite. Then there are $0_L, 1_L \in S$ satisfying*

$$A + B\begin{bmatrix} 0_L & 0_L & \dots & 0_L \\ 0_L & 0_L & \dots & 0_L \\ \vdots & \vdots & \ddots & \vdots \\ 0_L & 0_L & \dots & 0_L \end{bmatrix} = A\begin{bmatrix} 1_L & 0_L & \dots & 0_L \\ 0_L & 1_L & \dots & 0_L \\ \vdots & \vdots & \ddots & \vdots \\ 0_L & 0_L & \dots & 1_L \end{bmatrix} = A \tag{6.3}$$

and

$$A + \begin{bmatrix} 0_L & 0_L & \dots & 0_L \\ 0_L & 0_L & \dots & 0_L \\ \vdots & \vdots & \ddots & \vdots \\ 0_L & 0_L & \dots & 0_L \end{bmatrix} B = \begin{bmatrix} 1_L & 0_L & \dots & 0_L \\ 0_L & 1_L & \dots & 0_L \\ \vdots & \vdots & \ddots & \vdots \\ 0_L & 0_L & \dots & 1_L \end{bmatrix} A = A \tag{6.4}$$

for all $A, B \in L$.

Proof Take $K = \{A_{ij} : A \in L, 1 \leq i \leq m \text{ and } 1 \leq j \leq n\} \subseteq S$. Since L is non-empty and finite, and since each $A \in L$ has only finitely many entries, it is clear that K is a non-empty finite subset of S. Therefore, by Definition 4.1 (ii) for S, there are $0_K, 1_K \in S$ satisfying $a + b0_K = a1_K = a$ and $a + 0_K b = 1_K a = a$ for all $a, b \in K$.

Now let $A, B \in L$. Then repeated application of the fact that $a + b0_K = a$ for all $a, b \in K$ gives

$$\begin{aligned} A_{ij} + B_{i1}0_K + \cdots + B_{in}0_K &= A_{ij} + B_{i2}0_K + \cdots + B_{in}0_K \\ &= A_{ij} + B_{i3}0_K + \cdots + B_{in}0_K \\ &\vdots \\ &= A_{ij} + B_{in}0_K \\ &= A_{ij} \end{aligned} \tag{6.5}$$

for all $1 \leq i \leq m$ and all $1 \leq j \leq n$. Similarly

$$A_{ij}1_K + \sum_{\substack{k=1 \\ k\neq j}}^{n} A_{ij}0_K = A_{ij} + \sum_{\substack{k=1 \\ k\neq j}}^{n} A_{ij}0_K = A_{ij} \tag{6.6}$$

for all $1 \leq i \leq m$ and all $1 \leq j \leq n$, because we also have $a1_K = a$ for all $a \in K$, and thus (6.3) is satisfied with $0_L = 0_K$ and $1_L = 1_K$. A dual argument using the

fact that $a + 0_K b = 1_K a = a$ for all $a, b \in K$ shows that (6.4) is also satisfied with $0_L = 0_K$ and $1_L = 1_K$. □

An immediate consequence of Proposition 6.2 is that if $L \subseteq S^{n\times n}$ is non-empty and finite then there are matrices $0_L, I_L \in S^{n\times n}$ satisfying $A + B0_L = AI_L = A$ and $A + 0_L B = I_L A = A$ for all $A, B \in L$. Therefore Definition 4.1 (ii) holds for $S^{n\times n}$, and as such $S^{n\times n}$ is a semiring. When we specifically want to treat $S^{n\times n}$ as a semiring, rather than just as a set of matrices, we will refer to it as the *full matrix semiring* $\mathrm{M}_n(S)$. Subsemirings of $\mathrm{M}_n(S)$ will be called *matrix semirings*.

Another very useful consequence of Proposition 6.2 is that we can use matrix multiplication to "isolate" any desired column (or, dually, row) from a matrix. This ability is often taken for granted in the context of standard semirings because we can just multiply by a standard 'basis' vector to isolate a particular column (see Golan [34, Example 17.1]), but in the absence of global identity elements it is something less trivial that needs verifying. If $A \in S^{m\times n}$ then by Proposition 6.2 (with $L = \{A\}$) there is some $I_A \in S^{n\times n}$ satisfying $AI_A = A$, so if we write $v_1, \ldots, v_n \in S^{n\times 1}$ for the columns of I_A then the jth column of A is given by Av_j. Notice that each column vector v_j has jth entry 1_L and all other entries 0_L, by (6.3), and so the v_j are clearly analogues of the standard basis vectors we would normally use to isolate the columns of A. A variation on this argument gives us the following result.

Proposition 6.3 *Let S be a semiring and let $A, B \in S^{m\times n}$. If $Av = Bv$ for all $v \in S^{n\times 1}$ then $A = B$.*

Proof Take $L = \{A, B\}$. Then by Proposition 6.2 there is some $I_L \in S^{n\times n}$ satisfying $AI_L = A$ and $BI_L = B$. By the above argument, for each $1 \leq j \leq n$ the jth column of A is given by Av_j, where $v_j \in S^{n\times 1}$ denotes the jth column of I_L, and similarly the jth column of B is given by Bv_j. Therefore the jth columns of A and B are equal because, by the hypothesis, we have $Av_j = Bv_j$ for all $1 \leq j \leq n$. Hence $A = B$. □

Matrix multiplication is rather trivial in the case of 1×1 matrices, as we just recover multiplication on S (recall that we identified $S^{1\times 1}$ with S), but (6.2) permits 1×1 matrices to be multiplied by more than just other 1×1 matrices. Indeed, we can take any column vector $x \in S^{m\times 1}$ and multiply it by $a \in S$ to obtain another column vector $xa \in S^{m\times 1}$. This procedure has the effect of scaling (from the right) each entry of x by a. The fact that matrix multiplication is associative means that

we have $x(ab) = (xa)b$ for all $a, b \in S$ and all $x \in S^{m\times 1}$, and as such this scaling is a right action of S on $S^{m\times 1}$. Moreover, Definition 4.4 (i) holds for $S^{m\times 1}$ because $x(a+b) = xa + xb$ and $(x+y)a = xa + ya$ for all $a, b \in S$ and all $x, y \in S^{m\times 1}$, and thus $S^{m\times 1}$ is a right S-module provided it has right local identities. Proposition 6.2 immediately confirms the existence of right local identities, for if $L \subseteq S^{m\times 1}$ is non-empty and finite then there are $0_L, 1_L \in S$ satisfying $x + y0_L = x1_L = x$ for all $x, y \in L$, by (6.3). Therefore $S^{m\times 1}$ is a right S-module by Definition 4.4.

Note that we could also view $S^{m\times 1}$ as a left S-module, via the action which scales $x \in S^{m\times 1}$ on the left by $a \in S$, but it is more natural to view it as a right module because then the action coincides with matrix multiplication. So, unless S is commutative, we will treat $S^{m\times 1}$ as exclusively a right S-module. If we want a left S-module of vectors, a dual argument confirms that the natural choice is $S^{1\times m}$ with the action of S given by matrix multiplication on the other side.

If $A \in S^{m\times n}$ then the columns of A are vectors in $S^{m\times 1}$, and there are n of them, so taken together they are a non-empty finite subset of $S^{m\times 1}$. Given any such $L = \{x_1, \ldots, x_n\} \subseteq S^{m\times 1}$, a *right S-linear combination* of the x_j is an expression of the form $x_1a_1 + \cdots + x_na_n$ for $a_1, \ldots, a_n \in S$, and we write LS for the set of all right S-linear combinations of the vectors in L. If $x_1, \ldots, x_n \in L$ are given as the columns of a matrix $A \in S^{m\times n}$ then a right S-linear combination of the x_j is just an expression of the form Av for $v \in S^{n\times 1}$ (where the n entries of v are $a_1, \ldots, a_n$ above), and we write $\mathrm{Col}(A)$ instead of LS. Either way, $\mathrm{Col}(A) = LS$ is a right S-submodule of $S^{m\times 1}$ because it is closed under addition and the right action of S. The following definition records the construction of this very important module.

Definition 6.4 Let S be a semiring and let $A \in S^{m\times n}$. The *column space* of A is the right S-module

$$\mathrm{Col}(A) = \left\{Av : v \in S^{n\times 1}\right\} \subseteq S^{m\times 1} \tag{6.7}$$

of all right S-linear combinations of the columns of A.

Notice that $\mathrm{Col}(A)$ is precisely the image of the function $S^{n\times 1} \to S^{m\times 1}$ given by $v \mapsto Av$. This function is right S-linear, and is obviously surjective as a function $S^{n\times 1} \to \mathrm{Col}(A)$, but it need not be injective. In other words, two different vectors in $S^{n\times 1}$ could give rise to the same element of $\mathrm{Col}(A)$. In section 8 we will introduce a relation which records when this happens.

In addition to being a right S-module, the column space of a matrix $A \in S^{m\times n}$ is *finitely generated* because it comprises all linear combinations of a finite set of vectors. An important special case is when A only has one row ($m = 1$), so that its set L of "columns" is just a non-empty finite subset of S. In this case $\mathrm{Col}(A) = LS$ is a finitely generated right ideal of S, where a *right ideal* of S is a right S-submodule of S (i.e., a subset of S that is closed under addition and right multiplication by elements of S).

When we speak of modules generated by sets, we might have in mind a slightly different definition to the one given above: if L is a subset of a module X then we could define the module generated by L to be the smallest submodule of X that contains L, where 'smallest' means the intersection of all such submodules. The following result confirms that this definition is equivalent to the above definition for finitely generated submodules of $S^{m\times 1}$.

Proposition 6.5 *Let S be a semiring and let $A \in S^{m\times n}$. Then $\mathrm{Col}(A)$ is the intersection of all the right S-submodules of $S^{m\times 1}$ that contain the columns of A.*

Proof We begin by showing that $\mathrm{Col}(A)$ contains the columns of A. As described above, by Proposition 6.2 there is some $I_A \in S^{n\times n}$ satisfying $AI_A = A$. For each $1 \leq j \leq n$ the jth column of A is then given by Av_j, where v_j denotes the jth column of I_A, and thus each column of A is contained in $\mathrm{Col}(A)$ by (6.7). Therefore $\mathrm{Col}(A)$ is a right S-submodule of $S^{m\times 1}$ that contains the columns of A.

To show that $\mathrm{Col}(A)$ is the intersection of all such submodules, let X be a right S-submodule of $S^{m\times 1}$ and suppose that X contains the columns $x_1, \ldots, x_n \in S^{m\times 1}$ of A. Then since X is closed under addition and the right action of S, we have $x_1a_1 + \cdots + x_na_n \in X$ for all $a_1, \ldots, a_n \in S$. But, by definition, each element of $\mathrm{Col}(A)$ can be written as such a right S-linear combination for some $a_1, \ldots, a_n \in S$, and so $\mathrm{Col}(A) \subseteq X$. Hence $\mathrm{Col}(A)$ is the smallest right S-submodule of $S^{m\times 1}$ that contains the columns of A. □

The *row space* of a matrix $A \in S^{m\times n}$ is the finitely generated left S-submodule of $S^{1\times n}$ defined by

$$\mathrm{Row}(A) = \big\{uA : u \in S^{1\times m}\big\}. \tag{6.8}$$

That is, $\mathrm{Row}(A)$ comprises all *left S-linear combinations* of the rows of A, and, by Proposition 6.5 dual, $\mathrm{Row}(A)$ may also be characterised as the intersection of all the left S-submodules of $S^{1\times n}$ that contain the rows of A. In the special case

$n = 1$, the "rows" of A form a non-empty finite set $L \subseteq S$ and $\mathrm{Row}(A) = SL$ is a finitely generated *left ideal* of S (a subset of S that is closed under addition and left multiplication by elements of S).

To understand the multiplicative behaviour of matrices over a given semiring S, it is helpful (in the first instance) to attempt to describe the equivalence relations of Green [35] for matrices. Green's relations can be defined on any semigroup, so, in particular, can be defined on each semigroup $\left(S^{n\times n}, \cdot\right)$. However, as we will see below, Green's relations can actually be defined for matrices of any size without any extra difficulty; there is no great simplification in only considering square matrices of a fixed size. Hollings and Kambites [40, Proposition 3.1] have already considered Green's matrices over commutative 1-semirings, so since the outcome over a general semiring is essentially the same, we give only a brief summary of the key definitions and results.

Two matrices $A \in S^{m\times n}$ and $B \in S^{p\times q}$ are $\mathcal{R}$ *related*, written $A \mathcal{R} B$, if $m = p$ and there are $P \in S^{n\times q}$ and $Q \in S^{q\times n}$ satisfying $AP = B$ and $BQ = A$. In other words, $A \mathcal{R} B$ if and only if it is possible to get from A to B and back again by multiplication on the right. Dually, A and B are $\mathcal{L}$ *related*, written $A \mathcal{L} B$, if $n = q$ and there are $P \in S^{p\times m}$ and $Q \in S^{m\times p}$ satisfying $PA = B$ and $QB = A$.

Proposition 6.6 *Let S be a semiring, let $A \in S^{m\times n}$ and let $B \in S^{m\times q}$. Then $A \mathcal{R} B$ if and only if* $\mathrm{Col}(A) = \mathrm{Col}(B)$.

Proof By symmetry, it is sufficient to show that $AP = B$ for some $P \in S^{n\times q}$ if and only if $\mathrm{Col}(B) \subseteq \mathrm{Col}(A)$. If $AP = B$ for some $P \in S^{n\times q}$ then $\mathrm{Col}(B) \subseteq \mathrm{Col}(A)$ because $APv = Bv$ for all $v \in S^{q\times 1}$. Conversely, if $\mathrm{Col}(B) \subseteq \mathrm{Col}(A)$ then $\mathrm{Col}(A)$ contains each column of B by Proposition 6.5. This means that the jth column of B can be written as Av_j for some $v_j \in S^{n\times 1}$, and thus

$$B = \begin{bmatrix} Av_1 & \dots & Av_q \end{bmatrix} = A \begin{bmatrix} v_1 & \dots & v_q \end{bmatrix} \tag{6.9}$$

for some $\left[\begin{smallmatrix} v_1 & \cdots & v_q \end{smallmatrix}\right] \in S^{n\times q}$, as required. □

Green's $\mathcal{H}$ relation is defined to be the intersection of $\mathcal{L}$ and $\mathcal{R}$, that is, two matrices $A \in S^{m\times n}$ and $B \in S^{p\times q}$ are $\mathcal{H}$ *related* if and only if $A \mathcal{R} B$ and $A \mathcal{L} B$. Proposition 6.6 and its dual tell us that $A \mathcal{H} B$ if and only if $\mathrm{Col}(A) = \mathrm{Col}(B)$ and $\mathrm{Row}(A) = \mathrm{Row}(B)$. Finally, another relation that can be produced from $\mathcal{L}$ and $\mathcal{R}$ is their (relational) composition $\mathcal{D} = \mathcal{R} \circ \mathcal{L}$, which is defined as follows: A and B are

$\mathcal{D}$ *related* if there is some $C \in S^{m\times q}$ with $A \mathcal{R} C \mathcal{L} B$. It turns out that $A \mathcal{D} B$ if and only if there is some $D \in S^{p\times n}$ with $A \mathcal{L} D \mathcal{R} B$, and as such $\mathcal{D}$ is equal to the alternative composition $\mathcal{L} \circ \mathcal{R}$. It then follows that $\mathcal{D}$ is the smallest equivalence relation containing $\mathcal{L}$ and $\mathcal{R}$ (see Clifford and Preston [17, Lemma 2.1]).

Proposition 6.7 *Let S be a semiring, let $A \in S^{m\times n}$ and let $B \in S^{p\times q}$. If $A \mathcal{D} B$ then* $\mathrm{Col}(A) \cong \mathrm{Col}(B)$ *as right S-modules.*

Proof Since $A \mathcal{D} B$ there is some $C \in S^{m\times q}$ with $A \mathcal{R} C \mathcal{L} B$. This means that there are $P_A \in S^{n\times q}$ and $Q_A \in S^{q\times n}$ satisfying $AP_A = C$ and $CQ_A = A$, and also that there are $P_B \in S^{m\times p}$ and $Q_B \in S^{p\times m}$ satisfying $P_B B = C$ and $Q_B C = B$. Therefore

$$Q_B A v = Q_B C Q_A v = B Q_A v \in \mathrm{Col}(B) \tag{6.10}$$

for all $v \in S^{n\times 1}$, and so we can define a function $\mathrm{Col}(A) \to \mathrm{Col}(B)$ by $x \mapsto Q_B x$. Similarly, we can define a function $\mathrm{Col}(B) \to \mathrm{Col}(A)$ by $x \mapsto P_B x$ because

$$P_B B v = C v = A P_A v \in \mathrm{Col}(A) \tag{6.11}$$

for all $v \in S^{q\times 1}$. These functions are clearly right S-linear, so it remains to show that they are mutually inverse. This holds because we have

$$P_B Q_B A v = P_B B Q_A v = C Q_A v = A v \tag{6.12}$$

for all $v \in S^{n\times 1}$ and

$$Q_B P_B B v = Q_B C v = B v \tag{6.13}$$

for all $v \in S^{q\times 1}$. Hence $\mathrm{Col}(A) \cong \mathrm{Col}(B)$ as right S-modules. □

In section 9 we will explore a condition on S that allows us to characterise the $\mathcal{D}$ relation by isomorphism of column spaces (see Theorem 9.11).

7 Direct products and monoid semirings

In section 6 we saw a construction that takes a semiring S and produces another, related, semiring. Namely, if S is a semiring then we can form the full matrix semiring $\mathrm{M}_n(S)$ for each $n \in \mathbf{N}$. In this section we consider two more constructions that can be used to produce new semirings.

The first of these constructions takes two semirings S and T and makes the Cartesian product $S \times T$ a semiring. Addition and multiplication on $S \times T$ are defined componentwise, so that we have $(a_S, a_T) + (b_S, b_T) = (a_S + b_S, a_T + b_T)$ and $(a_S, a_T)(b_S, b_T) = (a_S b_S, a_T b_T)$ for all $(a_S, a_T), (b_S, b_T) \in S \times T$ (see Golan [34, page 19]). If $L \subseteq S \times T$ is non-empty and finite then the sets

$$L_S = \{a_S : (a_S, a_T) \in L\} \tag{7.1}$$

and

$$L_T = \{a_T : (a_S, a_T) \in L\} \tag{7.2}$$

are also non-empty and finite, so by Definition 4.1 (ii) for S and T there are local identities $0_{L_S}, 1_{L_S} \in S$ and $0_{L_T}, 1_{L_T} \in T$. Local identities for L are then given by $0_L = (0_{L_S}, 0_{L_T})$ and $1_L = (1_{L_S}, 1_{L_T})$, and with these local identities $S \times T$ is a semiring in the sense of Definition 4.1. We call this semiring the *direct product* of S and T.

The direct product construction is not limited to two semirings of course. If $S_1, \ldots, S_n$ are semirings then the n-fold product $S_1 \times \cdots \times S_n$ is also a semiring, so in particular the n-fold product $S^n = S \times \cdots \times S$ is a semiring for every semiring S.[1] Another construction related to S^n is that of a 'monoid semiring', defined below. The idea here is that addition is the same as on S^n, but multiplication is changed so as to be influenced by the structure of a fixed finite monoid.

Definition 7.1 Let S be a semiring and let $(M, \cdot, 1)$ be a finite monoid. The *monoid semiring* SM is the set of functions $M \to S$ with addition and multiplication given by

$$(V + W)r = Vr + Wr \tag{7.3}$$

and

$$(VW)r = \sum_{r=st} Vs \cdot Wt \tag{7.4}$$

respectively for all $V, W \in SM$ and all $r \in M$. Note that '$r = st$' here means "all $s, t \in M$ satisfying $r = st$".

[1] We can actually consider arbitrary direct products of semirings, as even if there are infinitely many semirings, a non-empty finite subset of the Cartesian product has only finitely many elements in each component. In other words, infinite direct products of semirings still have local identities.

Multiplication on a monoid semiring is often called 'convolution' (see Golan [34, page 29]). Note that it is sometimes possible to consider monoid semirings for infinite monoids. For instance, if M is the free monoid on a finite set then the sum in (7.4) is always finite, or if $S = \mathbf{B}$ then the sum in (7.4) exists regardless of whether it is finite. Alternatively, if S is a 0-semiring then we could take all finitely supported functions $M \to S$ instead of all functions $M \to S$ (see Golan [32, Chapter 4]).

Even in the case of a finite monoid M though, we have not made sure that SM is always a semiring in the sense of Definition 4.1. To do this, we need to find local identities for each non-empty finite $L \subseteq SM$. Since M is finite, the image of each $V \in L$ is a non-empty finite subset of S, so we can define a non-empty finite $K \subseteq S$ by $K = \{Vr : V \in L \text{ and } r \in M\}$. By Definition 4.1 (ii) for S there are local identities $0_K, 1_K \in S$, and an argument similar to the proof of Proposition 6.2 then shows that the functions $0_L, 1_L \in SM$ given by

$$0_L r = 0_K \tag{7.5}$$

and

$$1_L r = \begin{cases} 1_K & \text{if } r = 1, \\ 0_K & \text{otherwise} \end{cases} \tag{7.6}$$

for all $r \in M$ are local identities for L.

An element of a Boolean monoid semiring $\mathbf{B}M$ assigns either a 0 or a 1 to each element of M, so instead of working with functions $V \in \mathbf{B}M$ we can just work with subsets $\{s \in M : Vs = 1\}$ of M. In other words, we can identify $\mathbf{B}M$ with the powerset of M. After this identification, the sum and product of subsets $V, W \in \mathbf{B}M$ are given by $V \cup W$ and $\{st : s \in V \text{ and } t \in W\}$ respectively, because addition and multiplication on $\mathbf{B}$ are given by disjunction and conjunction respectively. The fact that $\mathbf{B}$ is a standard semiring means that it is not necessary to consider local identities; $\mathbf{B}M$ is a standard semiring, with additive identity element $\emptyset$ and multiplicative identity element $\{1\} \subseteq M$. As mentioned above, there is nothing about this construction that fails if M is infinite, and thus there is a standard semiring $\mathbf{B}M$ for every monoid M. However, not all of our results concerning Boolean monoid semirings apply in the case M is infinite (see Theorem 20.4).

If $G = \{r_1, \ldots, r_n\}$ is a finite group of order n then we call SG a *group semiring*.

In this case the product of $V, W \in SG$ is given by

$$(VW)r = \sum_{j=1}^{n} Vr_j \cdot W\big(r_j^{-1}r\big) = \sum_{i=1}^{n} V\big(rr_i^{-1}\big) \cdot Wr_i \tag{7.7}$$

for all $r \in G$ because $r = st$ for $s, t \in G$ if and only if $t = s^{-1}r$, which also happens if and only if $s = rt^{-1}$. The expressions in (7.7) are reminiscent of the definition of matrix multiplication, so we might expect SG to have an interpretation as a matrix semiring. Indeed, the following result confirms this suspicion.

Lemma 7.2 *Let S be a semiring and let $G = \{r_1, \ldots, r_n\}$ be a finite group of order n. Then the function $f\colon SG \to \mathrm{M}_n(S)$ given by $(fV)_{ij} = V\big(r_i^{-1}r_j\big)$ for all $1 \leq i, j \leq n$ is an injective homomorphism.*

Proof If $V \in SG$ then the first row of fV is

$$\begin{bmatrix} V\big(r_1^{-1}r_1\big) & \ldots & V\big(r_1^{-1}r_n\big) \end{bmatrix}, \tag{7.8}$$

so if we assume (without loss of generality) that $r_1 = 1$ then we have $Vr_j = (fV)_{1j}$ for all $1 \leq j \leq n$. Therefore V can be recovered from the first row of fV, and thus f is injective.

To show that f is a homomorphism, we first need to verify Definition 5.3 (i). That is, we need to check that $f(V+W) = fV + fW$ for all $V, W \in SG$. This follows immediately from (7.3) and (6.1). Next we need to check that $f(VW) = (fV)(fW)$ for all $V, W \in SG$. By (7.7) we have

$$(f(VW))_{ik} = (VW)\big(r_i^{-1}r_k\big) = \sum_{j=1}^{n} Vr_j \cdot W\big(r_j^{-1}r_i^{-1}r_k\big) \tag{7.9}$$

for all $1 \leq i, k \leq n$, which we can rewrite as

$$(f(VW))_{ik} = \sum_{j=1}^{n} V\big(r_i^{-1}r_j\big) \cdot W\big(r_j^{-1}r_k\big) = \sum_{j=1}^{n} (fV)_{ij} \cdot (fV)_{jk} \tag{7.10}$$

because $\big(r_i^{-1}r_j\big)^{-1} r_i^{-1}r_k = r_j^{-1}r_k$. Therefore $f(VW) = (fV)(fW)$ by (6.2), and as such Definition 5.3 (ii) holds for f.

Finally, we need to verify Definition 5.3 (iii). That is, for each non-empty finite $L \subseteq \mathrm{M}_n(S)$ we need to find $0_L, I_L \in \mathrm{Im}(f)$ satisfying $A + B0_L = AI_L = A$ and

$A + 0_L B = I_L A = A$ for all $A, B \in L$. Proposition 6.2 confirms the existence of local identities $0_L, I_L \in \mathrm{M}_n(S)$, so if we can show that in fact $0_L, I_L \in \mathrm{Im}(f)$ then we will be done. Notice that $\mathrm{Im}(f)$ contains each matrix of the form

$$\begin{bmatrix} a & b & \dots & b \\ b & a & \dots & b \\ \vdots & \vdots & \ddots & \vdots \\ b & b & \dots & a \end{bmatrix} \tag{7.11}$$

for $a, b \in S$, because this matrix is the image of the element of SG given by $r_1 \mapsto a$ and $r_i \mapsto b$ for all $1 < i \leq n$. In particular then, $\mathrm{Im}(f)$ contains the local identities $0_L, I_L \in \mathrm{M}_n(S)$ of Proposition 6.2, and consequently Definition 5.3 (iii) holds for f. Hence f is a homomorphism. □

Lemma 7.2 tells us that if S is a semiring and G is a finite group of order n then we can view the group semiring SG as a subsemiring of the full matrix semiring $\mathrm{M}_n(S)$. It also allows us to obtain the following powerful result, which implies that the properties of SG are closely related to the properties of $\mathrm{M}_n(S)$. We will make use of this fact in Corollary 12.5.

Theorem 7.3 *Let S be a semiring and let G be a finite group of order n. Then*

(i) *$SG \cong S^{1\times n}$ as right SG-modules;*

(ii) *$(SG)^{n\times 1} \cong \mathrm{M}_n(S)$ as right SG-modules; and*

(iii) *SG is a right retract of $\mathrm{M}_n(S)$.*

Proof As above, write $G = \{r_1, \dots, r_n\}$ with $r_1 = 1$ and define $f\colon SG \to \mathrm{M}_n(S)$ by $(fV)_{ij} = V\big(r_i^{-1} r_j\big)$ for all $1 \leq i, j \leq n$. Then f is an injective homomorphism by Lemma 7.2. Note that f is right SG-linear because the right action of SG on $\mathrm{M}_n(S)$ is defined via f (see Proposition 5.6).

(i). Write g for the function $\mathrm{M}_n(S) \to S^{1\times n}$ that selects the first row of each matrix. It is clear that g is right $\mathrm{M}_n(S)$-linear, where the right action of $\mathrm{M}_n(S)$ on $S^{1\times n}$ is given by matrix multiplication, so in particular g is right $\mathrm{Im}(f)$-linear. This means that g is right SG-linear, because, by Proposition 5.6, the right actions of SG on $\mathrm{M}_n(S)$ and $S^{1\times n}$ are defined via f. Therefore $g \circ f\colon SG \to S^{1\times n}$ is right

SG-linear. If $V \in SG$ then

$$(g \circ f)V = \begin{bmatrix} (fV)_{11} & \dots & (fV)_{1n} \end{bmatrix} = \begin{bmatrix} Vr_1 & \dots & Vr_n \end{bmatrix}, \tag{7.12}$$

so $g \circ f$ is a bijection, and as such $g \circ f$ is an isomorphism of right SG-modules.

(ii). This follows immediately from (i), as $g \circ f \colon SG \to S^{1\times n}$ can be extended to a right SG-module isomorphism $(SG)^{1\times n} \to \mathrm{M}_n(S)$ by applying it entrywise.

(iii). Since $f \colon SG \to \mathrm{M}_n(S)$ is an injective homomorphism we can identify SG with $\mathrm{Im}(f)$, and so to show that SG is a right retract of $\mathrm{M}_n(S)$ we need to construct a right SG-linear function $h \colon \mathrm{M}_n(S) \to SG$ satisfying $(h \circ f)V = V$ for all $V \in SG$ (see Definition 5.7). Take $h = (g \circ f)^{-1} \circ g$. Then h is right SG-linear with $h \circ f = (g \circ f)^{-1} \circ (g \circ f)$, and as such $h \circ f$ is the identity function on SG. Hence SG is a right retract of $\mathrm{M}_n(S)$. □

The three main problems

8 Kernels and separation

Our aim is to present a systematic way to understand linear algebra over semirings, so we now describe three problems whose solutions give useful information about the behaviour of matrices with entries in an arbitrary semiring. In section 9 we study the problem of extending linear functionals defined on row and column spaces, and in section 10 we propose a way to explain the observed variety of relationships (e.g., isomorphism, conjugate isomorphism and anti-isomorphism) between row and column spaces. We begin in this section by introducing two kinds of 'kernel'.

Our first kind of kernel is the kernel of a set of vectors. As we show in Proposition 8.6 (i), the kernel of the row space of a matrix A records when two column vectors become equal in the column space of A. In other words, the kernel of $\mathrm{Row}(A)$ measures how far away the function that takes a column vector into $\mathrm{Col}(A)$ is from being a bijection (see Definition 6.4). Given a semiring S, our first main problem is to describe the kernel classes of the row space (and the column space) of an arbitrary matrix with entries in S. As we will see in sections 13 and 18, the techniques used to describe the structure of kernel classes must be tailored to the semiring in question, and as such it is difficult to give a general approach to this problem.

Our other kind of kernel is the kernel of a relation on vectors, or, more usefully, the 'double kernel' of a set of vectors. The double kernel is a closure operator that enlarges a set of vectors by including all other vectors that behave in the same way with respect to multiplication, and so by studying the double kernel of the row space of a matrix A we can find out whether matrix multiplication tells us everything there is to know about A. In section 9 we will link the double kernel of $\mathrm{Row}(A)$ with linear functions on $\mathrm{Col}(A)$ and show that if the double kernel is strictly bigger than $\mathrm{Row}(A)$ then there are more linear functions on $\mathrm{Col}(A)$ than matrix multiplication alone can account for.

Definition 8.1 Let S be a semiring and let $X \subseteq S^{1\times n}$. The *kernel* of X is the relation

$$\operatorname{Ker}(X) = \big\{(v, v') \in S^{n\times 1} \times S^{n\times 1} : xv = xv' \text{ for all } x \in X\big\}. \tag{8.1}$$

It is clear from the form of (8.1) that if $X \subseteq S^{1\times n}$ then $\operatorname{Ker}(X)$ is an equivalence relation on $S^{n\times 1}$. Moreover, the fact that (8.1) is defined in terms of multiplication means that $\operatorname{Ker}(X)$ is actually a right S-congruence on $S^{1\times n}$. This allows us to treat the set $S^{n\times 1}/\operatorname{Ker}(X)$ of equivalence classes as a right S-module (see page 30), but in view of Proposition 8.6 (ii), below, there is really no need to study this particular module in the case X is the row space of a matrix. However, for completeness, we at least show that $\operatorname{Ker}(X)$ is always a congruence.

Proposition 8.2 *Let S be a semiring and let $X \subseteq S^{1\times n}$. Then $\operatorname{Ker}(X)$ is a right S-congruence on $S^{n\times 1}$.*

Proof To show that $\operatorname{Ker}(X)$ is a right S-congruence on $S^{n\times 1}$ we first need to show that it is compatible with addition, so let $(v, v'), (w, w') \in \operatorname{Ker}(X)$. Then $xv = xv'$ and $xw = xw'$ for all $x \in X$, from which we have $x(v+w) = xv+xw = xv'+xw' = x(v'+w')$ for all $x \in X$. Therefore $(v+w, v'+w') \in \operatorname{Ker}(X)$, and as such $\operatorname{Ker}(X)$ is compatible with addition.

We also need to show that $\operatorname{Ker}(X)$ is compatible with the right action, so let $a \in S$ and let $(v, v') \in \operatorname{Ker}(X)$. Then $xv = xv'$ for all $x \in X$, which means that $x(va) = (xv)a = (xv')a = x(v'a)$ for all $x \in X$. Therefore $(va, v'a) \in \operatorname{Ker}(X)$, and as such $\operatorname{Ker}(X)$ is compatible with the right action of S. Hence $\operatorname{Ker}(X)$ is a right S-congruence on $S^{1\times n}$. □

Notice that Definition 8.1 does not restrict $X \subseteq S^{1\times n}$ in any way, i.e., X does not need to be a right S-submodule of $S^{1\times n}$. This means that we can apply $\operatorname{Ker}(-)$ to the empty subset of $S^{1\times n}$, vacuously giving $\operatorname{Ker}(\emptyset) = S^{n\times 1} \times S^{n\times 1}$. At the other extreme, we have $\operatorname{Ker}\big(S^{1\times n}\big) = \big\{(v, v) : v \in S^{n\times 1}\big\}$ because if $v, v' \in S^{n\times 1}$ with $xv = xv'$ for all $x \in S^{1\times n}$ then $v = v'$ by Proposition 6.3 dual. Proposition 8.5, below, lists some more properties of $\operatorname{Ker}(-)$ as a function from subsets of $S^{1\times n}$ to relations on $S^{n\times 1}$.

The kernel of $X \subseteq S^{1\times n}$ comprises all pairs of column vectors that give the same product with each element of X, but once we have these pairs we could immediately reverse the process and consider the set of row vectors in $S^{1\times n}$ that give the same

product with each element of $\mathrm{Ker}(X)$. We write $\mathrm{Ker}^2(X)$ for this set, and we call it the *double kernel* of X. To make this notation and terminology more meaningful, we also introduce the following definition.

Definition 8.3 Let S be a semiring and let $F \subseteq S^{n\times 1} \times S^{n\times 1}$. The *kernel* of F is the set

$$\mathrm{Ker}(F) = \big\{y \in S^{1\times n} : yv = yv' \text{ for all } (v, v') \in F\big\} \tag{8.2}$$
$$= \big\{y \in S^{1\times n} : F \subseteq \mathrm{Ker}(y)\big\}. \tag{8.3}$$

We have now defined two different notions of kernel: one which takes a set of row vectors to a relation on column vectors, and one which takes a relation on column vectors back to a set of row vectors. In particular, we are now able to meaningfully take a "double" kernel by simply computing $\mathrm{Ker}\,\mathrm{Ker}(X)$ for $X \subseteq S^{1\times n}$, and it is clear from Definition 8.3 that this construction agrees with the double kernel $\mathrm{Ker}^2(X)$ described above. That is, we have

$$\mathrm{Ker}^2(X) = \mathrm{Ker}\,\mathrm{Ker}(X) = \big\{y \in S^{1\times n} : \mathrm{Ker}(X) \subseteq \mathrm{Ker}(y)\big\}, \tag{8.4}$$

where $\mathrm{Ker}(y)$ is shorthand for $\mathrm{Ker}(\{y\})$. The relationship between our two notions of kernel is captured by the following result.

Proposition 8.4 *If S is a semiring then*

$$X \subseteq \mathrm{Ker}(F) \quad \Leftrightarrow \quad F \subseteq \mathrm{Ker}(X) \tag{8.5}$$

for all $X \subseteq S^{1\times n}$ and all $F \subseteq S^{n\times 1} \times S^{n\times 1}$.

Proof Let $X \subseteq S^{1\times n}$ and let $F \subseteq S^{n\times 1} \times S^{n\times 1}$. By (8.3) we have $X \subseteq \mathrm{Ker}(F)$ if and only if $F \subseteq \mathrm{Ker}(x)$ for all $x \in X$, and thus it is sufficient to show that $F \subseteq \mathrm{Ker}(X)$ if and only if $F \subseteq \mathrm{Ker}(x)$ for all $x \in X$. But this is clear because we have

$$\mathrm{Ker}(X) = \bigcap_{x\in X} \mathrm{Ker}(x) \tag{8.6}$$

by (8.1). □

Another way to express Proposition 8.4 is to say that taking kernels constitutes a Galois connection (see Definition 15.1) between subsets of $S^{1\times n}$ and relations on

$S^{n\times 1}$. Specifically, we have two posets, namely $\mathrm{Pow}\big(S^{1\times n}\big)$ and $\mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big)$ ordered by inclusion, and two functions

$$\mathrm{Pow}\big(S^{n\times 1}\big) \underset{\mathrm{Ker}(-)}{\overset{\mathrm{Ker}(-)}{\rightleftarrows}} \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big) \tag{8.7}$$

satisfying (8.5) for all $X \in \mathrm{Pow}\big(S^{1\times n}\big)$ and all $F \in \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big)$. We then automatically obtain the following properties of $\mathrm{Ker}(-)$ as a function from subsets of $S^{1\times n}$ to relations on $S^{n\times 1}$, and vice versa.

Proposition 8.5 *If S is a semiring then*

(i) *the functions*

$$\mathrm{Ker}(-)\colon\ \mathrm{Pow}\big(S^{1\times n}\big) \to \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big) \tag{8.8}$$

and

$$\mathrm{Ker}(-)\colon\ \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big) \to \mathrm{Pow}\big(S^{1\times n}\big) \tag{8.9}$$

are inclusion-reversing;

(ii) *the functions*

$$\mathrm{Ker}^2(-)\colon\ \mathrm{Pow}\big(S^{1\times n}\big) \to \mathrm{Pow}\big(S^{1\times n}\big) \tag{8.10}$$

and

$$\mathrm{Ker}^2(-)\colon\ \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big) \to \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big) \tag{8.11}$$

are inclusion-preserving and expanding; and

(iii) *we have $\mathrm{Ker}^3(X) = \mathrm{Ker}(X)$ for all $X \in \mathrm{Pow}\big(S^{1\times n}\big)$ and $\mathrm{Ker}^3(F) = \mathrm{Ker}(F)$ for all $F \in \mathrm{Pow}\big(S^{n\times 1}\times S^{n\times 1}\big)$.*

Proof Apply Proposition 15.2 to Proposition 8.4. □

Viewed as a function from subsets of $S^{1\times n}$ to relations on $S^{n\times 1}$, the fact that $\mathrm{Ker}^2(-)$ is expanding means that we have $X \subseteq \mathrm{Ker}^2(X)$ for all $X \subseteq S^{1\times n}$. We can also see from (8.4) and (8.1) that $\mathrm{Ker}^2(X)$ comprises all $y \in S^{1\times n}$ with the property that $yv = yv'$ for all $v, v' \in S^{n\times 1}$ satisfying $xv = xv'$ for all $x \in X$, so $\mathrm{Ker}^2(X)$ enlarges X by including all $y \in S^{1\times n}$ that—as far as multiplication is

concerned—look just like elements of X. The elements of $\mathrm{Ker}^2(X)$ will therefore be called *inseparable* from X, while the elements of $S^{1\times n} \setminus \mathrm{Ker}^2(X)$ will be called *separable* from X. These definitions are in the spirit of 'separation' in the context of topological vector spaces (see Aliprantis and Border [2, Chapter 5]).

Proposition 8.5 (iii) tells us that $\mathrm{Ker}^4(-) = \mathrm{Ker}^2(-)$, and as such $\mathrm{Ker}^2(-)$ is idempotent. This makes $\mathrm{Ker}^2(-)$ a closure operator on $S^{1\times n}$ (see page 96), and means that there are no vectors inseparable from $\mathrm{Ker}^2(X)$ that were originally separable from $X \subseteq S^{1\times n}$. It therefore makes sense to think of $\mathrm{Ker}^2(X)$ as being a better behaved "completion" of X, and so if we want linear algebra over S to be well-behaved then we should insist that $\mathrm{Ker}^2(-)$ does not strictly enlarge row spaces of matrices with entries in S. We will discuss this condition in section 9.

Now recall that our first main problem is to describe the equivalence classes of $\mathrm{Ker}\,\mathrm{Row}(A)$ for an arbitrary matrix $A \in S^{m\times n}$. As the following result demonstrates, these special kernels are easier to understand than the general kernels discussed above.

Proposition 8.6 *Let S be a semiring and let $A \in S^{m\times n}$. Then*

(i) $\mathrm{Ker}\,\mathrm{Row}(A) = \big\{(v, v') \in S^{n\times 1} \times S^{n\times 1} : Av = Av'\big\}$; *and*

(ii) $S^{n\times 1}/\,\mathrm{Ker}\,\mathrm{Row}(A) \cong \mathrm{Col}(A)$ *as right S-modules.*

Proof (i). Let $v, v' \in S^{n\times 1}$. If $(v, v') \in \mathrm{Ker}\,\mathrm{Row}(A)$ then we have $uAv = uAv'$ for all $u \in S^{1\times m}$ by (8.1), and thus $Av = Av'$ by Proposition 6.3 dual. Conversely, if $Av = Av'$ then $uAv = uAv'$ for all $u \in S^{1\times m}$, and as such $(v, v') \in \mathrm{Ker}\,\mathrm{Row}(A)$ by (8.1) again. Hence $\mathrm{Ker}\,\mathrm{Row}(A) = \big\{(v, v') \in S^{n\times 1} \times S^{n\times 1} : Av = Av'\big\}$.

(ii). It is clear from (i) that $\mathrm{Ker}\,\mathrm{Row}(A)$ is just the (set-theoretic) kernel of the surjective right S-linear function $S^{n\times 1} \to \mathrm{Col}(A)$ given by $v \mapsto Av$. Therefore $S^{n\times 1}/\,\mathrm{Ker}\,\mathrm{Row}(A) \cong \mathrm{Col}(A)$ as right S-modules by Proposition 5.2. □

Proposition 8.6 (i) tells us that if $A \in S^{m\times n}$ then $\mathrm{Ker}\,\mathrm{Row}(A)$ records the pairs of column vectors that get identified in $\mathrm{Col}(A)$, but this does not reveal anything about the "internal" structure of the equivalence classes of $\mathrm{Ker}\,\mathrm{Row}(A)$. The "external" structure of the classes is described by Proposition 8.6 (ii), however, and so if we understood the relationship between the row and column spaces of A (section 10 deals with this problem) then we would be able to completely describe the relationship between $\mathrm{Row}(A)$ and the external structure of its kernel.

We are not able to say anything general about the internal structure of the classes of $\operatorname{Ker}\operatorname{Row}(A)$ because their contents depends on what type of semiring S is. One possible way to describe $\operatorname{Ker}\operatorname{Row}(A)$ in the case of a particular semiring would be to find a smaller relation $F \subseteq S^{n\times 1} \times S^{n\times 1}$ that captures the essential information about each equivalence class of $\operatorname{Ker}\operatorname{Row}(A)$. Specifically, we could look for a small, easy to understand relation F satisfying $\operatorname{Ker}\operatorname{Row}(A) = \operatorname{Ker}^2(F)$, and then explain how to enlarge F to get $\operatorname{Ker}^2(F)$. We will take this approach in section 18.

9 Extensions and exactness

Linear functions are of fundamental importance in linear algebra, but until now we have been exclusively focussing on matrices and vectors. We begin this section by considering linear functions as objects in their own right, and we show that if A is a matrix with entries in a semiring S then the set of all right S-linear functions $\operatorname{Col}(A) \to S$ is a left S-module. There are two reasons why we only prove this particular restricted result. Firstly, we are mainly interested in the behaviour of matrices anyway, and secondly, the way we have set up S-modules means that a general result along the lines of "if X is a right S-module then the set of right S-linear functions $X \to S$ is a left S-module" is not even possible. If we tried to prove such a result we would run into technical difficulties establishing the existence of local identities, but, as we will see, working with column spaces means that there is ultimately something finite (the matrix) that we can use to produce local identities for linear functions.

To help understand linear functions on column spaces, we consider when such a function can be extended to a linear function on all appropriately-sized column vectors—the most desirable situation being, of course, when each function has an extension. This motivates our central definition: a semiring is (right) 'exact' if each linear function on the column space of a matrix has an extension (see Definition 9.2).

Our next main problem after describing kernels (see section 8) is to decide whether a given semiring is exact, or to at least provide some useful information about linear functions, so it would be convenient if exactness was linked with kernels in some way. As we show in Lemma 9.9, it turns out that there is a correspondence between the linear functions on the column space of a matrix A and the elements of the double kernel of the row space of A, and this leads to a characterisation of

exactness in terms of double kernels.

Finally, we discuss some immediate consequences of exactness and we propose some possible ways to show that particular semirings are exact. We will consider many more consequences of exactness in sections 11 and 12.

Proposition 9.1 *Let S be a semiring and let $A \in S^{m\times n}$. Then the set*

$$\mathrm{Col}(A)^* = \{f\colon \mathrm{Col}(A) \to S : f \text{ is right } S\text{-linear}\} \tag{9.1}$$

is a left S-module.

Proof To show that $\mathrm{Col}(A)^*$ is a left S-module, we first need to define addition and a left action of S. The sum of $f, g \in \mathrm{Col}(A)^*$ is given by $(f+g)x = fx + gx$ for all $x \in \mathrm{Col}(A)$ and the left action of $a \in S$ on $f \in \mathrm{Col}(A)$ is given by $(af)x = a(fx)$ for all $x \in \mathrm{Col}(A)$. It is easy to see that $f+g\colon \mathrm{Col}(A) \to S$ and $af\colon \mathrm{Col}(A) \to S$ will be right S-linear, so they are indeed elements of $\mathrm{Col}(A)^*$. It is also clear that $(ab)f = a(bf)$ for all $a, b \in S$ and all $f \in \mathrm{Col}(A)^*$, and as such we really have defined an action of S on $\mathrm{Col}(A)^*$ (see page 18).

Next we need to verify Definition 4.4 (i) dual for $\mathrm{Col}(A)^*$, that is, we need to check that $(a+b)f = af + bf$ and $a(f+g) = af + ag$ for all $a, b \in S$ and all $f, g \in \mathrm{Col}(A)^*$. The proof of this is routine, and in fact no part of the proof so far relies on any particular property of $\mathrm{Col}(A)$; this portion of the result holds for any right S-module.

To verify Definition 4.4 (ii) dual for $\mathrm{Col}(A)^*$, we need to establish the existence of *left local identities* $0_L, 1_L \in S$ satisfying $f + 0_L g = 1_L f = f$ for all $f, g \in L$ whenever $L \subseteq \mathrm{Col}(A)^*$ is non-empty and finite. Given such a set L, we define a non-empty finite set $K \subseteq S^{1\times n}$ by $K = \{fA : f \in L\}$, where $fA \in S^{1\times n}$ denotes the row vector obtained by applying f to the columns of A. This construction is described in more detail on page 58, but the important point is that fA satisfies $f(Av) = (fA)v$ for all $v \in S^{n\times 1}$. Now let $f, g \in L$. By Proposition 6.2 there are $0_K, 1_K \in S$ satisfying

$$fA + 0_K(gA) = 1_K(fA) = fA, \tag{9.2}$$

and thus

$$f(Av) + 0_K(g(Av)) = 1_K(f(Av)) = f(Av) \tag{9.3}$$

for all $v \in S^{n\times 1}$. The above definitions of addition and the left action of S on

$\mathrm{Col}(A)^*$ then give

$$(f + 0_K g)x = (1_K f)x = fx \tag{9.4}$$

for all $x \in \mathrm{Col}(A)$, which means that $f + 0_K g = 1_K f = f$. Therefore L has left local identities $0_L = 0_K$ and $1_L = 1_K$, and as such Definition 4.4 (ii) dual is satisfied for $\mathrm{Col}(A)^*$. Hence $\mathrm{Col}(A)^*$ is a left S-module. □

The left S-module $\mathrm{Col}(A)^*$ is called the *dual* of $\mathrm{Col}(A)$, and its elements are often called *linear functionals* on $\mathrm{Col}(A)$. As mentioned above, we are interested in when a function $f \in \mathrm{Col}(A)^*$ can be extended, so we should define precisely what we mean by this. If S is a semiring and $A \in S^{m\times n}$ then an *extension* of $f \in \mathrm{Col}(A)^*$ is a right S-linear function $g\colon S^{m\times 1} \to S$ satisfying $gx = fx$ for all $x \in \mathrm{Col}(A)$. The dual of $\mathrm{Row}(A)$ and the notion of an extension of a linear functional $f \in \mathrm{Row}(A)^*$ are defined dually.

Definition 9.2 A semiring S is

(i) *right exact* if each $f \in \mathrm{Col}(A)^*$ has an extension whenever $A \in S^{m\times n}$; and is

(ii) *left exact* if each $f \in \mathrm{Row}(A)^*$ has an extension whenever $A \in S^{m\times n}$.

A semiring which is both left and right exact will be called *exact* (see Wilding et al. [86, section 3]). The most obvious examples of exact semirings are fields, for if R is a field and $A \in R^{m\times n}$ then each $f \in \mathrm{Col}(A)^*$ has an extension by standard results of linear algebra (see Roman [72, Theorem 1.4 and page 103]). We will give more examples of exact semirings in sections 14 and 19.

Example 9.3 The ring $\mathbf{Z}$ is not exact. For instance, if the function $f \in (2\mathbf{Z})^*$ given by $2b \mapsto b$ for all $b \in \mathbf{Z}$ had an extension g then, by linearity, we would have $g1 = 1/2$, which is impossible.

There is nothing special about $\mathbf{Z}$ in Example 9.3; the same trick shows that any integral domain that is not a field cannot be exact. An *integral domain* is a commutative non-trivial ring R in which each $a \in R \setminus \{0\}$ is not a *zero divisor*, that is, there is no $b \in R \setminus \{0\}$ with $ab = 0$ (see Cohn [21, page 117]). If $a \neq 0$ is any non-invertible element of R then the function $f \in (aR)^*$ given by $ab \mapsto b$ for all $b \in R$ is well-defined (because a is not a zero divisor) and does not have an extension (because a is not invertible). Therefore R cannot be exact unless it is a field.

The above argument is essentially a reformulation of a result of Wilding et al. [86, Proposition 4.2], which says that in an exact commutative ring each non-zero element is a zero divisor or is invertible. Proposition 13.5 generalises this result to matrices.

The condition for a semiring S to be right exact is very similar to the condition for S to be injective as a right S-module (see Golan [34, page 197]). In fact, as we will see below, exactness is just a restricted form of such 'self-injectivity' of S. Semirings which are known to be self-injective include the Boolean semiring (see Wang [83, Lemma 1]), proper quotients of principal ideal domains (see Lam [56, Example 3.12]) and certain Boolean rings (see Lam [56, Corollary 3.11D]).

Definition 9.4 A semiring S is *right self-injective* if whenever X and Y are right S-modules and there are right S-linear functions $f\colon X \to S$ and $h\colon X \to Y$ with h injective, there is some right S-linear function $g\colon Y \to S$ satisfying $g \circ h = f$.

Definition 9.4 is much easier to comprehend (and remember) in the form of a diagram: a semiring S is right self-injective if and only if whenever $f\colon X \to S$ is a right S-linear function and $h\colon X \to Y$ is an injective right S-linear function, there is some right S-linear function $g\colon Y \to S$ that makes the diagram

$$\begin{array}{ccc} X & \xrightarrow{\ h\ } & Y \\ {\scriptstyle f}\downarrow & \swarrow{\scriptstyle g} & \\ S & & \end{array} \tag{9.5}$$

commute. In the case of exactness we have the same basic diagram, but h is only permitted to be the inclusion of a column space into its containing module. That is, S is right exact if and only if whenever $A \in S^{m\times n}$ and $f\colon \mathrm{Col}(A) \to S$ is a right S-linear function, there is some right S-linear function $g\colon S^{m\times 1} \to S$ that makes the diagram

$$\begin{array}{ccc} \mathrm{Col}(A) & \hookrightarrow & S^{m\times 1} \\ {\scriptstyle f}\downarrow & \swarrow{\scriptstyle g} & \\ S & & \end{array} \tag{9.6}$$

commute. This proves the following result.

Proposition 9.5 *If S is a right self-injective semiring then S is right exact.*

Proof If S is right self-injective then S is right exact because, as described above, (9.6) is an instance of (9.5). □

Exactness for rings is usually called *FP-injectivity* (see Nicholson and Yousif [70, page 95]) because, as we have just seen, it is a restricted form of self-injectivity.[1] An FP-injective ring is also sometimes called 'absolutely pure', after Maddox [65]. We will briefly discuss the connection between FP-injectivity and self-injectivity for rings in section 12.

Although exactness of a semiring follows from self-injectivity, we will not actually show that any particular semiring is exact by establishing that it is self-injective. For one thing, an exact semiring need not be self-injective (see page 89), but the main reason we will not use self-injectivity to deduce exactness is that this approach does not give us enough information about the behaviour of matrices. In fact, we will not even show that any particular semiring is exact by directly checking Definition 9.2. Instead, we will characterise exactness in terms of double kernels and subsequently obtain exactness as a by-product of our attempt to understand kernels.

The first step towards casting exactness in terms of double kernels is to replace linear functionals by vectors. If S is a semiring and $f \in \mathrm{Col}(A)^*$ for some $A \in S^{m\times n}$ then since $\mathrm{Col}(A)$ contains the columns of A, we can apply f to each column in turn to obtain n elements of S. These n elements can then be written as a row vector (in the same order as the columns from which they came). We write $fA \in S^{1\times n}$ for this vector, and it is straightforward to verify that right S-linearity of f implies that we have $f(Av) = (fA)v$ for all $v \in S^{n\times 1}$. Alternatively, we could take this to be the defining property of fA.

The row vector fA contains all the information necessary to compute a linear function $f \in \mathrm{Col}(A)^*$ because if we want to find fx for a given $x \in \mathrm{Col}(A)$ then, by the above property, we have $fx = f(Av) = (fA)v$ for some $v \in S^{n\times 1}$. To actually do this in practice, though, we would have to work backwards from x and construct some $v \in S^{n\times 1}$ satisfying $Av = x$. This is not necessarily an easy problem, and, by Proposition 8.6 (i), is related to the problem of describing $\operatorname{Ker}\mathrm{Row}(A)$, so fA is not necessarily the most useful vector encapsulating f.

[1] FP stands for 'finitely presented'.

It would be better if there was some $u \in S^{1\times m}$ with $fx = ux$ for all $x \in \mathrm{Col}(A)$, then we could directly compute f by simply multiplying two vectors, but such a vector u does not necessarily exist. It turns out that there is some $u \in S^{1\times m}$ with this property if and only if f has an extension.

Proposition 9.6 *Let S be a semiring, let $A \in S^{m\times n}$ and let $f \in \mathrm{Col}(A)^*$. Then f has an extension if and only if there is some $u \in S^{1\times m}$ satisfying $fx = ux$ for all $x \in \mathrm{Col}(A)$.*

Proof ($\Rightarrow$). Suppose that f has an extension $g\colon S^{m\times 1} \to S$. By Proposition 6.2 there is some $I_A \in S^{m\times m}$ with $I_A A = A$, so since g is defined on the whole of $S^{m\times 1}$ we can apply it to the columns of I_A and obtain a vector $gI_A \in S^{1\times m}$. Therefore $g(Av) = g(I_A Av) = (gI_A)Av$ for all $v \in S^{n\times 1}$ because g is right S-linear, and as such $fx = gx = (gI_A)x$ for all $x \in \mathrm{Col}(A)$ because g is an extension of f.

($\Leftarrow$). Suppose that there is some $u \in S^{1\times m}$ with $fx = ux$ for all $x \in \mathrm{Col}(A)$. Then the right S-linear function $S^{m\times 1} \to S$ given by $y \mapsto uy$ is clearly an extension of f. □

Note that Proposition 9.6 certainly does not claim that every extension of f can be written in the form $y \mapsto uy$ for some $u \in S^{1\times m}$. It merely tells us that if f has an extension then there is a (possibly different) extension that can be written in this form. This issue is not something we need to worry about for standard semirings, as the existence of a global identity matrix I then ensures that any extension g of f can be written as $y \mapsto (gI)y$, but, as the following example illustrates, when S only has local identities there may be extensions of f that cannot be written like this.

Example 9.7 Let $A \in \mathbf{FT}^{2\times 2}$ and define a function $f\colon \mathrm{Col}(A) \to \mathbf{FT}$ by $x \mapsto x_{11}$, where $\mathbf{FT} = (\mathbf{R}, \max, +)$ denotes the finitary tropical semiring. Since f just selects the first entry of each vector in $\mathrm{Col}(A)$, it is clear that f is right $\mathbf{FT}$-linear (our choice of $\mathbf{FT}$ is not significant yet), and as such $f \in \mathrm{Col}(A)^*$. Moreover, there is no reason to restrict the definition of f to the vectors in $\mathrm{Col}(A)$; the function $\mathbf{FT}^{2\times 1} \to \mathbf{FT}$ given by $y \mapsto y_{11}$ is obviously an extension of f.

Suppose that this extension of f can be written as $y \mapsto uy$ for some $u \in \mathbf{FT}^{1\times 2}$. Then we have

$$y_{11} = uy = \max\{u_{11} + y_{11}, u_{12} + y_{21}\} \tag{9.7}$$

for all $y \in \mathbf{FT}^{2\times 1}$. In the case $y_{21} = 0$ this means that $u_{12} \leq y_{11}$ for all $y_{11} \in \mathbf{FT}$, which says that $\mathbf{FT}$ has a bottom element u_{12}. This is a contradiction, because $\mathbf{FT}$

has no bottom element, and so we conclude that f has an extension that cannot be written in the form $y \mapsto uy$ for any $u \in \mathbf{FT}^{1\times 2}$.

Proposition 9.8 *Let S be a semiring, let $A \in S^{m\times n}$ and let $f \in \mathrm{Col}(A)^*$. Then f has an extension if and only if $fA \in \mathrm{Row}(A)$.*

Proof By Proposition 9.6, f has an extension if and only if there is some $u \in S^{1\times m}$ satisfying $fx = ux$, which happens if and only if there is some $u \in S^{1\times m}$ with $(fA)v = f(Av) = uAv$ for all $v \in S^{1\times n}$. Therefore, by Proposition 6.3, f has an extension if and only if there is some $u \in S^{1\times m}$ with $fA = uA$, that is, if and only if $fA \in \mathrm{Row}(A)$. □

While we might not necessarily have $fA \in \mathrm{Row}(A)$, we do at least always have $fA \in \mathrm{Ker}^2 \mathrm{Row}(A)$. This is because if $v, v' \in S^{n\times 1}$ with $Av = Av'$ then

$$(fA)v = f(Av) = f(Av') = (fA)v' \tag{9.8}$$

because f is right S-linear. In other words, $\mathrm{Ker}\,\mathrm{Row}(A) \subseteq \mathrm{Ker}(fA)$, which means that $fA \in \mathrm{Ker}^2 \mathrm{Row}(A)$ by Definition 8.3. Conversely, each element of $\mathrm{Ker}^2 \mathrm{Row}(A)$ gives rise to an element of $\mathrm{Col}(A)^*$, and this allows us to show that $\mathrm{Col}(A)^*$ and $\mathrm{Ker}^2 \mathrm{Row}(A)$ are in fact isomorphic as left S-modules. Note that $\mathrm{Ker}^2 \mathrm{Row}(A)$ is a left S-submodule of $S^{1\times n}$ because we have $\mathrm{Ker}(y) \cap \mathrm{Ker}(z) \subseteq \mathrm{Ker}(y+z)$ and $\mathrm{Ker}(y) \subseteq \mathrm{Ker}(ay)$ for all $a \in S$ and all $y, z \in S^{1\times n}$.

Lemma 9.9 *Let S be a semiring and let $A \in S^{m\times n}$. Then the function from $\mathrm{Col}(A)^*$ to $\mathrm{Ker}^2 \mathrm{Row}(A)$ given by $f \mapsto fA$ is an isomorphism of left S-modules, with inverse given by $y \mapsto (Av \mapsto yv)$.*

Proof The function given by $f \mapsto fA$ is left S-linear because of how we made $\mathrm{Col}(A)^*$ a left S-module (see the proof of Proposition 9.1), so it remains to show that the proposed inverse is correct.

Let $y \in \mathrm{Ker}^2 \mathrm{Row}(A)$. Then $\mathrm{Ker}\,\mathrm{Row}(A) \subseteq \mathrm{Ker}(y)$ by Definition 8.3, which means that if $v, v' \in S^{n\times 1}$ with $Av = Av'$ then $yv = yv'$, and as such the function given by $Av \mapsto yv$ is well-defined. Call this function f. It is clear that f is right S-linear, so $f \in \mathrm{Col}(A)^*$, and since we have $(fA)v = (fAv) = yv$ for all $v \in S^{n\times 1}$ it follows immediately from Proposition 6.3 that $fA = y$. Therefore the composition $\mathrm{Ker}^2 \mathrm{Row}(A) \to \mathrm{Col}(A)^* \to \mathrm{Ker}^2 \mathrm{Row}(A)$ is the identity function.

Now let $f \in \mathrm{Col}(A)^*$ so that $fA \in \mathrm{Ker}^2 \mathrm{Row}(A)$, as described above. Then the function given by $Av \mapsto (fA)v$ is just f because $(fA)v = f(Av)$ for all $v \in S^{n\times 1}$, and as such the composition $\mathrm{Col}(A)^* \to \mathrm{Ker}^2 \mathrm{Row}(A) \to \mathrm{Col}(A)^*$ is also the identity function. Hence $\mathrm{Col}(A)^* \cong \mathrm{Ker}^2 \mathrm{Row}(A)$ as left S-modules. □

An immediate consequence of Lemma 9.9 is that if S is a semiring and $A \in S^{m\times n}$ then $\mathrm{Ker}^2 \mathrm{Row}(A) = \{fA : f \in \mathrm{Col}(A)^*\}$. Together with the above results, this fact gives us our working characterisation of exactness.

Proposition 9.10 *Let S be a semiring. Then S is right exact if and only if $\mathrm{Ker}^2 \mathrm{Row}(A) = \mathrm{Row}(A)$ for all $A \in S^{m\times n}$.*

Proof ($\Rightarrow$). Suppose that S is right exact and let $A \in S^{m\times n}$. We already know that $\mathrm{Row}(A) \subseteq \mathrm{Ker}^2 \mathrm{Row}(A)$, by Proposition 8.5 (ii), so it is sufficient to show that the reverse inclusion also holds. Now let $y \in \mathrm{Ker}^2 \mathrm{Row}(A)$. Then $y = fA$ for some $f \in \mathrm{Col}(A)^*$, by the above observation. Since S is right exact, f has an extension by Definition 9.2 (i), and thus $y = fA \in \mathrm{Row}(A)$ by Proposition 9.8. Hence $\mathrm{Ker}^2 \mathrm{Row}(A) = \mathrm{Row}(A)$.

($\Leftarrow$). Suppose that $\mathrm{Ker}^2 \mathrm{Row}(A) = \mathrm{Row}(A)$ for all $A \in S^{m\times n}$. To show that S is right exact we need to verify that each $f \in \mathrm{Col}(A)^*$ has an extension whenever $A \in S^{m\times n}$, so let $A \in S^{m\times n}$ and let $f \in \mathrm{Col}(A)^*$. Then $fA \in \mathrm{Ker}^2 \mathrm{Row}(A)$ by the above observation, and so $fA \in \mathrm{Row}(A)$, which means that f has an extension by Proposition 9.8 again. Hence S is right exact by Definition 9.2 (i). □

Proposition 9.10 tells us that exactness can be phrased in terms of a separation property: a semiring S is right exact if and only if each $y \in S^{1\times n} \setminus \mathrm{Row}(A)$ is separable from $\mathrm{Row}(A)$ whenever $A \in S^{m\times n}$. This is the way that Hollings and Kambites [40, Lemma 4.1] originally formulated exactness, and it allowed them to show that the finitary tropical semiring $\mathbf{FT} = (\mathbf{R}, \max, +)$ and the completed tropical semiring $\overline{\mathbf{T}} = (\mathbf{R}\cup\{-\infty, \infty\}, \max, +)$ are both exact. They were also able to obtain Theorem 9.11 (ii), below, for the tropical semiring $\mathbf{T} = (\mathbf{R}\cup\{-\infty\}, \max, +)$, even though it is an open question whether $\mathbf{T}$ is exact.

Theorem 9.11 *Let S be a right exact semiring, let $A \in S^{m\times n}$ and let $B \in S^{p\times q}$. Then*

(i) $\mathrm{Col}(A)^* \cong \mathrm{Row}(A)$ *as left S-modules; and*

(ii) $\mathrm{Col}(A) \cong \mathrm{Col}(B)$ *as right* S*-modules if and only if* $A \mathcal{D} B$.

Proof (i). By Lemma 9.9, $\mathrm{Col}(A)^* \cong \mathrm{Ker}^2 \mathrm{Row}(A)$ as left S-modules, so if S is right exact then $\mathrm{Col}(A)^* \cong \mathrm{Row}(A)$ as left S-modules by Proposition 9.10.

(ii). If $A \mathcal{D} B$ then $\mathrm{Col}(A) \cong \mathrm{Col}(B)$ as right S-modules by Proposition 6.7, so it remains to show that $A \mathcal{D} B$ whenever $\mathrm{Col}(A) \cong \mathrm{Col}(B)$.

Let $f\colon \mathrm{Col}(A) \to \mathrm{Col}(B)$ be an isomorphism of right S-modules and for each $1 \leq i \leq p$ write f_i for the right S-linear function $\mathrm{Col}(A) \to S$ given by $x \mapsto (fx)_{i1}$. Then by Proposition 9.6 there are $u_1, \ldots, u_p \in S^{1\times m}$ with $(fx)_{i1} = f_i x = u_i x$ for all $x \in \mathrm{Col}(A)$ and all $1 \leq i \leq p$. Therefore f is given by

$$fx = \begin{bmatrix} u_1 x \\ \vdots \\ u_p x \end{bmatrix} = \begin{bmatrix} u_1 \\ \vdots \\ u_p \end{bmatrix} x \tag{9.9}$$

for all $x \in \mathrm{Col}(A)$, and as such there is a matrix $P \in S^{p\times m}$ with $fx = Px$ for all $x \in \mathrm{Col}(A)$. A similar argument shows that there is a matrix $Q \in S^{m\times p}$ with $f^{-1}x = Qx$ for all $x \in \mathrm{Col}(B)$.

The fact that $f^{-1} \circ f$ is the identity function on $\mathrm{Col}(A)$ means that $QPAv = Av$ for all $v \in S^{n\times 1}$, and thus $QPA = A$ by Proposition 6.3. Therefore $A \mathcal{L} PA$. Also, since the image of f is $\mathrm{Col}(B)$ we have

$$\mathrm{Col}(B) = \{Px : x \in \mathrm{Col}(A)\} = \left\{PAv : v \in S^{n\times 1}\right\} = \mathrm{Col}(PA), \tag{9.10}$$

and as such $B \mathcal{R} PA$ by Proposition 6.6. Hence $A \mathcal{D} B$ because $A \mathcal{L} PA \mathcal{R} B$. □

Recall from section 8 that to understand $\mathrm{Ker}\,\mathrm{Row}(A)$ for a given $A \in S^{m\times n}$, it would be helpful if we could find a smaller relation $F \subseteq S^{n\times 1} \times S^{n\times 1}$ satisfying $\mathrm{Ker}^2(F) = \mathrm{Ker}\,\mathrm{Row}(A)$. The idea is that such a relation would be easier to describe, yet would still capture the essential information about $\mathrm{Ker}\,\mathrm{Row}(A)$. One way for F to satisfy $\mathrm{Ker}^2(F) = \mathrm{Ker}\,\mathrm{Row}(A)$ would of course be if $\mathrm{Ker}(F) = \mathrm{Row}(A)$, and it turns out that if each $A \in S^{m\times n}$ has such a relation F then we obtain right exactness of S in addition to a useful description of $\mathrm{Ker}\,\mathrm{Row}(A)$.

Lemma 9.12 *Let* S *be a semiring. Then* S *is right exact if for each* $A \in S^{m\times n}$ *there is some* $F \subseteq S^{n\times 1} \times S^{n\times 1}$ *satisfying* $\mathrm{Ker}(F) = \mathrm{Row}(A)$.

Proof Let $A \in S^{m\times n}$ and suppose that there is some $F \subseteq S^{n\times 1} \times S^{n\times 1}$ satisfying $\mathrm{Ker}(F) = \mathrm{Row}(A)$. Then $\mathrm{Ker}^3(F) = \mathrm{Ker}^2 \mathrm{Row}(A)$, where $\mathrm{Ker}^3(F) = \mathrm{Ker}(F)$ by Proposition 8.5 (iii), and thus we have $\mathrm{Ker}^2 \mathrm{Row}(A) = \mathrm{Row}(A)$. Hence S is right exact by Proposition 9.10. □

We will apply Lemma 9.12 in the proof of Theorem 19.4 to show that certain tropical-like semirings are exact.

10 Conjugations on semirings

Our last main problem is to explain the relationship between the row space and column space of a matrix with entries in a given semiring. It is not clear how best to approach this problem, as, in general, the row space of a matrix A could be nothing like the column space of A. On the other hand, $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ could be isomorphic, as in the case of a field, and so we should not expect to be able to say anything meaningful about the relationship between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ unless the semiring in question has some particular structure that we can make use of. To simplify matters, we will restrict our attention to cases where there is an isomorphism-like connection between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ for each matrix A.

Motivated by the fact that the row space and column space of a matrix with entries in $\mathbf{C}$ are conjugate isomorphic as well as isomorphic, we introduce a notion of a 'conjugation' on a semiring S. Our definition of a conjugation is formulated so as to induce a bijection between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ for each $A \in S^{m\times n}$. Moreover, these bijections will be structure preserving in a way that depends on the particular conjugation chosen. We will give more examples of conjugations and their induced "conjugate isomorphisms" in Theorems 14.2 and 19.5.

Since a conjugation should ultimately connect two objects (namely the row space and column space of a matrix), we will require a conjugation on a semiring S to be an involution on S. Note that by an *involution* on S we just mean a function $\overline{}\colon S \to S$ satisfying $\overline{\overline{a}} = a$ for all $a \in S$; we do not require $\overline{}$ to interact with addition or multiplication in any way. If $\overline{}$ also satisfies $\overline{a+b} = \overline{a} + \overline{b}$ and $\overline{ab} = \overline{b}\,\overline{a}$ for all $a, b \in S$ then we will call $\overline{}$ a *standard involution* (see Golan [33, page 68]).

Given an involution $\overline{}$ on a semiring S, we extend $\overline{}$ to matrices with entries in S via the transpose operation. That is, we define $\overline{A} \in S^{n\times m}$ by $\overline{A}_{ji} = \overline{A_{ij}}$ for each $A \in S^{m\times n}$, so that we have $\overline{A+B} = \overline{A} + \overline{B}$ and $\overline{AB} = \overline{B}\,\overline{A}$ for all $A, B \in S^{m\times n}$

if $\overline{}$ happens to be a standard involution. For instance, if S is commutative then the identity function on S is a standard involution, and thus when this involution is extended to matrices we simply recover the familiar properties $(A+B)^{\mathrm{T}} = A^{\mathrm{T}} + B^{\mathrm{T}}$ and $(AB)^{\mathrm{T}} = B^{\mathrm{T}}A^{\mathrm{T}}$. We can now properly define what we mean by a conjugation on S.

Definition 10.1 Let S be a semiring and let $\overline{}$ be an involution on S. Then $\overline{}$ is a *conjugation* on S if for each $A \in S^{m\times n}$ there are $M, N \in S^{m\times n}$ satisfying

(i) $M\overline{uA} = A\overline{uM}$ and $\overline{M\overline{uA}}N = uA$ for all $u \in S^{1\times m}$; and

(ii) $\overline{Av}N = \overline{Nv}A$ and $M\overline{\overline{Av}N} = Av$ for all $v \in S^{n\times 1}$.

Definition 10.1 looks impenetrable, but the idea behind it is quite simple. For each $A \in S^{m\times n}$ a conjugation should induce a bijection between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$, so we insist that there are functions $\mathrm{Row}(A) \to \mathrm{Col}(A)$ and $\mathrm{Col}(A) \to \mathrm{Row}(A)$ that somehow rely on conjugation. The two functions we have in mind are

$$\mathrm{Row}(A) \xrightarrow{x \mapsto M\overline{x}} \mathrm{Col}(A) \xrightarrow{x \mapsto \overline{x}N} \mathrm{Row}(A), \tag{10.1}$$

and all that (i) and (ii) are saying is that these functions have the correct codomains and are mutually inverse. Specifically, $M\overline{uA} = A\overline{uM}$ and $\overline{Av}N = \overline{Nv}A$ (which are actually a little stronger than is strictly necessary) ensure that $M\overline{x} \in \mathrm{Col}(A)$ for all $x \in \mathrm{Row}(A)$ and that $\overline{x}N \in \mathrm{Row}(A)$ for all $x \in \mathrm{Col}(A)$, while $\overline{M\overline{uA}}N = uA$ and $M\overline{\overline{Av}N} = Av$ ensure that $x \mapsto M\overline{x}$ and $x \mapsto \overline{x}N$ compose to give the identity functions on $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ respectively.

The bijection $x \mapsto M\overline{x}$ that a conjugation induces between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ is a linear function (multiplication by M) composed with the conjugation itself, so it is reasonable to expect the interaction of this bijection with the S-module structures of $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ to reflect the way the conjugation interacts with addition and multiplication on S. In particular, if $\overline{}$ is a standard involution then, as the following result shows, the induced bijection between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ is very similar to an isomorphism of S-modules.

Proposition 10.2 *Let S be a semiring, let $\overline{}$ be a standard involution on S and let $A \in S^{m\times n}$. If $\overline{}$ is a conjugation then there is a bijection $f\colon \mathrm{Row}(A) \to \mathrm{Col}(A)$ satisfying $f(x+y) = fx + fy$ and $f(ax) = (fx)\overline{a}$ for all $a \in S$ and all $x, y \in \mathrm{Row}(A)$.*

Proof Write f for the bijection $x \mapsto M\overline{x}$ described above. Then since $\overline{\ }$ is a semiring involution we have

$$f(x+y) = M\overline{x+y} = M(\overline{x}+\overline{y}) = M\overline{x} + M\overline{y} = fx + fy \tag{10.2}$$

and

$$f(ax) = M\overline{ax} = M\overline{x}\,\overline{a} = (fx)\overline{a} \tag{10.3}$$

for all $a \in S$ and all $x, y \in \mathrm{Row}(A)$. □

In cases where S is commutative and the identity function on S is a conjugation (see Theorem 14.2), Proposition 10.2 tells us that $\mathrm{Row}(A) \cong \mathrm{Col}(A)$ as S-modules for all $A \in S^{m\times n}$. As we show in Theorem 10.4, below, complex conjugation on $\mathbf{C}$ is a conjugation in the sense of Definition 10.1, and so in this case Proposition 10.2 recovers the fact that there is a conjugate isomorphism between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ for all $A \in \mathbf{C}^{m\times n}$ (see Roman [72, Theorem 9.18]). To help show that complex conjugation is a conjugation in our sense, and to make it easier to show that other standard involutions are conjugations, we first prove the following useful result.

Lemma 10.3 *Let S be a semiring and let $\overline{\ }$ be a standard involution on S. Then $\overline{\ }$ is a conjugation if for each $A \in S^{n\times n}$ there are $M, N \in S^{n\times n}$ satisfying $M\overline{A} = A\overline{M}$ and $A\overline{M}N = A$.*

Proof Let $B \in S^{m\times n}$. We need to find $M_B, N_B \in S^{m\times n}$ satisfying Definition 10.1 (i) and (ii) for B, but we cannot directly use the hypothesis unless B is square. So suppose that $m \geq n$ and write $A \in S^{m\times m}$ for the square matrix obtained by repeating some columns of B as necessary. It is obvious that A and B have the same column space, and thus $A \,\mathcal{R}\, B$ by Proposition 6.6. This means that there are $P \in S^{m\times n}$ and $Q \in S^{n\times m}$ with $AP = B$ and $BQ = A$. By the hypothesis, there are also $M_A, N_A \in S^{n\times n}$ with $M_A\overline{A} = A\overline{M_A}$ and $A\overline{M_A}N_A = A$.

Now take $M_B = M_A\overline{Q}$ and $N_B = N_AP$. To satisfy Definition 10.1 (i) for B we first need to show that $M_B\overline{uB} = B\overline{uM_B}$ for all $u \in S^{1\times m}$, but since $\overline{\ }$ is a standard involution it is sufficient to show that $M_B\overline{B} = B\overline{M_B}$. This holds because we have

$$M_B\overline{B} = M_A\overline{Q}\,\overline{B} = M_A\overline{A} = A\overline{M_A} = BQ\overline{M_A} = B\overline{M_B}. \tag{10.4}$$

Next we need to show that $\overline{M_B\overline{uB}}N_B = uB$ for all $u \in S^{1\times m}$. For this it is sufficient

to show that $B\overline{M_B}N_B = B$, and indeed we have

$$B\overline{M_B}N_B = BQ\overline{M_A}N_AP = A\overline{M_A}N_AP = AP = B. \tag{10.5}$$

To satisfy Definition 10.1 (ii) for B it is sufficient to show that $\overline{B}N_B = \overline{N_B}B$ and $M_B\overline{N_B}B = B$. For the first of these we have

$$\overline{B}N_B = \overline{N_B}M_B\overline{B}N_B = \overline{N_B}B\overline{M_B}N_B = \overline{N_B}B \tag{10.6}$$

by (10.4) and (10.5), and for the second we have

$$M_B\overline{N_B}B = M_B\overline{B}N_B = B\overline{M_B}N_B = B \tag{10.7}$$

by (10.4), (10.5) and (10.6).

By repeating some rows of B instead of some columns, a dual argument confirms that Definition 10.1 (i) and (ii) are also satisfied for B in the case $m \leq n$. Hence $\overline{\ }$ is a conjugation. □

Theorem 10.4 *Complex conjugation* $\overline{\ }: \mathbf{C} \to \mathbf{C}$ *is a conjugation in the sense of Definition 10.1.*

Proof Let $A \in \mathbf{C}^{n\times n}$. As a consequence of the singular value decomposition of A (see Roman [72, page 445]) there are $P, Q \in \mathbf{C}^{n\times n}$ with P and Q unitary, that is, with $\overline{P} = P^{-1}$ and $\overline{Q} = Q^{-1}$, and with PAQ real diagonal. Therefore $PAQ = \overline{PAQ}$, because taking the conjugate transpose of a (square) real diagonal matrix has no effect, and thus

$$\overline{P}\,\overline{Q}\,\overline{A} = \overline{P}\big(\overline{Q}\,\overline{A}\,\overline{P}\big)P = \overline{P}(PAQ)P = AQP \tag{10.8}$$

because $\overline{\ }$ is a standard involution. We also have

$$AQP\overline{P}\,\overline{Q} = AQ\overline{Q} = A, \tag{10.9}$$

and as such A satisfies the hypothesis of Lemma 10.3 with $M = N = \overline{P}\,\overline{Q}$. Hence $\overline{\ }$ is a conjugation. □

Transferring exactness

11 Exact matrix and product semirings

Before giving examples of exact semirings (see sections 14 and 19), we would like to further develop the abstract theory of exactness so that we can understand the examples in context. Specifically, we are interested in the extent to which exactness can be transferred. If a semiring S is exact, what can be said about subsemirings of S, semirings containing S as a subsemiring and semirings constructed from S?

In this section we show that the full matrix and direct product constructions described in sections 6 and 7 (respectively) preserve exactness. When combined with the general results about subsemirings in section 12, these facts allow us to easily show that more semirings related to S are exact whenever S is. For instance, if S is exact then we can deduce that the group semiring SG is exact for every finite group G by realising SG as a matrix semiring and applying exactness of full matrix semirings (see Corollary 12.5).

Since we are considering the problem of transferring exactness from one semiring to another, it will be necessary to refer to row spaces and kernels in more than one semiring at a time. To remove any ambiguity, we will write $\mathrm{Row}_S(A)$ and $\mathrm{Ker}_S \mathrm{Row}_S(A)$ for the row space of a matrix A and its kernel when A is intended to be interpreted as a matrix with entries in S. This distinction is particularly important in our first result, as a matrix with entries in $\mathrm{M}_n(S)$ can also be regarded as a block matrix with entries in S.

Proposition 11.1 *If S is a right exact semiring then each $\mathrm{M}_n(S)$ is right exact.*

Proof For convenience, we write T instead of $\mathrm{M}_n(S)$ in this proof.

Let $A \in T^{p\times q}$ and let $y \in \mathrm{Ker}^2_T \mathrm{Row}_T(A)$. Each entry of A is an $n \times n$ matrix with entries in S, so we can treat A as a $pn \times qn$ matrix with entries in S. That is, we identity $T^{p\times q}$ and $S^{pn\times qn}$, and similarly we have $y \in T^{1\times q} = S^{n\times qn}$ with rows

$y_1, \ldots, y_n \in S^{1\times qn}$.

We want to show that each $y_i \in \mathrm{Ker}_S^2 \mathrm{Row}_S(A)$, so let $v, v' \in S^{qn\times 1}$ with $Av = Av'$. Now write $w, w' \in S^{qn\times n} = T^{q\times 1}$ for the matrices obtained by repeating v and v', so that $w = [\, v \cdots v \,]$ and $w' = [\, v' \cdots v' \,]$, then we have $Aw = Aw'$ over T because $Av = Av'$ over S. Therefore $(w, w') \in \mathrm{Ker}_T \mathrm{Row}_T(A)$, and thus $(w, w') \in \mathrm{Ker}_T(y)$ because $y \in \mathrm{Ker}_T^2 \mathrm{Row}_T(A)$. This means that $yw = yw'$ over T, so $y_i v = y_i v'$ over S for all $1 \leq i \leq n$ because w and w' were defined using v and v' respectively. Hence each $y_i \in \mathrm{Ker}_S^2 \mathrm{Row}_S(A)$ because, as we have just shown, $\mathrm{Ker}_S \mathrm{Row}_S(A) \subseteq \mathrm{Ker}_S(y_i)$ for all $1 \leq i \leq n$.

We can now use right exactness of S. By Proposition 9.10 each $y_i \in \mathrm{Row}_S(A)$, so for each $1 \leq i \leq n$ there is some $u_i \in S^{1\times pn}$ with $y_i = u_i A$. Combining these vectors, we obtain a matrix

$$u = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix} \in S^{n\times pn} = T^{1\times p} \tag{11.1}$$

which satisfies $y = uA$ over T, and as such $y \in \mathrm{Row}_T(A)$. We have therefore shown that $\mathrm{Ker}_T^2 \mathrm{Row}_T(A) \subseteq \mathrm{Row}_T(A)$, and so in fact $\mathrm{Ker}_T^2 \mathrm{Row}_T(A) = \mathrm{Row}_T(A)$. Hence T is right exact by Proposition 9.10. □

Our other useful result in this section is that the direct product of exact semirings is an exact semiring. This result is not at all surprising, and is quite easy to prove, but unfortunately the proof is rather off-putting. The fundamental observation is that we can split any matrix over $S \times T$ into a matrix over S (the "left part") and a matrix over T (the "right part") using the obvious correspondence between $(S \times T)^{m\times n}$ and $S^{m\times n} \times T^{m\times n}$.

Proposition 11.2 *If S and T are right exact semirings then $S \times T$ is right exact.*

Proof Let $A \in (S \times T)^{m\times n}$ and let $y \in \mathrm{Ker}_{S\times T}^2 \mathrm{Row}_{S\times T}(A)$. Each entry of A is a pair, so we write $A_S \in S^{m\times n}$ for the matrix comprising the left-hand entries of A and $A_T \in T^{m\times n}$ for the matrix comprising the right-hand entries of A. Similarly, we write $y_S \in S^{1\times n}$ and $y_T \in T^{1\times n}$ for the left and right parts of y.

We want to show that $y_S \in \mathrm{Ker}_S^2 \mathrm{Row}_S(A_S)$ and $y_T \in \mathrm{Ker}_T^2 \mathrm{Row}_T(A_T)$, so let $v_S, v'_S \in S^{n\times 1}$ with $A_S v_S = A_S v'_S$ and let $v_T, v'_T \in T^{n\times 1}$ with $A_T v_T = A_T v'_T$. Then

$Av = Av'$, where $v, v' \in (S \times T)^{n\times 1}$ denote the vectors with left parts v_S, v'_S respectively and right parts v_T, v'_T respectively. Therefore $(v, v') \in \operatorname{Ker}_{S\times T} \operatorname{Row}_{S\times T}(A)$, and thus $(v, v') \in \operatorname{Ker}_{S\times T}(y)$ because $y \in \operatorname{Ker}^2_{S\times T} \operatorname{Row}_{S\times T}(A)$. This means that $yv = yv'$, with left part $y_S v_S = y_S v'_S$ and right part $y_T v_T = y_T v'_T$, and as such

$$\operatorname{Ker}_S \operatorname{Row}_S(A_S) \subseteq \operatorname{Ker}_S \operatorname{Row}_S(y_S) \tag{11.2}$$

and

$$\operatorname{Ker}_T \operatorname{Row}_T(A_T) \subseteq \operatorname{Ker}_T \operatorname{Row}_T(y_T). \tag{11.3}$$

Hence $y_S \in \operatorname{Ker}^2_S \operatorname{Row}_S(A_S)$ and $y_T \in \operatorname{Ker}^2_T \operatorname{Row}_T(A_T)$.

We are now ready to use right exactness of S and T. By Proposition 9.10 we have $y_S \in \operatorname{Row}_S(A_S)$ and $y_T \in \operatorname{Row}_T(A_T)$, so there are $u_S \in S^{1\times m}$ and $u_T \in S^{1\times n}$ with $y_S = u_S A_S$ and $y_T = u_T A_T$. Therefore $y = uA$, where $u \in (S \times T)^{1\times m}$ denotes the vector with left part u_S and right part u_T, and as such $y \in \operatorname{Row}_{S\times T}(A)$. To summarise, we have shown that $\operatorname{Ker}^2_{S\times T} \operatorname{Row}_{S\times T}(A) \subseteq \operatorname{Row}_{S\times T}(A)$, and so in fact $\operatorname{Ker}^2_{S\times T} \operatorname{Row}_{S\times T}(A) = \operatorname{Row}_{S\times T}(A)$. Hence Proposition 9.10 confirms that $S \times T$ is right exact. □

12 Ideals and exact subsemirings

In this section we consider what effect exactness of a semiring S has on the subsemirings of S and on the semirings containing S as a subsemiring. We begin by showing that if S is exact then it must be contained in every larger semiring in a "rigid" way. This places quite a strong restriction on the finitely generated ideals of any semiring that contains S as a subsemiring, and allows us to show, for instance, that an exact integral domain must be a field.

The significance of finitely generated ideals then leads us to wonder whether "partial" exactness of a semiring S (i.e., at the level of finitely generated ideals) always transfers to genuine exactness of S. In other words, when studying exactness, is it sufficient to just consider the finitely generated ideals of S? This question is also motivated by the well-known result of Baer [5] which characterises self-injectivity of a ring purely in terms of its ideals.

Finally, we ask when exactness can be transferred down from a semiring to a subsemiring. Our main result (Theorem 12.4) gives sufficient conditions for a subsemiring of an exact semiring to be exact, and is of most use in the case of a full

matrix semiring. In section 11 we showed that each $\mathrm{M}_n(S)$ is exact if S is an exact semiring, so the results in this section allow us to transfer exactness of S to various matrix semirings. In particular, as we deduce in Corollary 12.5, exactness of S can be transferred to the group semiring SG for every finite group G.

Note that, as in section 11, we use a subscript to indicate the semiring in which the entries of a matrix lie.

Lemma 12.1 *Let T be a semiring and let S be a subsemiring of T. If S is right exact then*

(i) $\mathrm{Row}_T(A) \cap S^{1\times n} = \mathrm{Row}_S(A)$*; and*

(ii) $\mathrm{Row}_S(A) \subseteq \mathrm{Row}_S(B)$ *if and only if* $\mathrm{Row}_T(A) \subseteq \mathrm{Row}_T(B)$

for all $A \in S^{m\times n}$ and all $B \in S^{p\times n}$.

Proof (i). It is clear that $\mathrm{Row}_S(A) \subseteq \mathrm{Row}_T(A) \cap S^{1\times n}$, so it remains to show that the reverse inclusion also holds. Let $x \in \mathrm{Row}_T(A) \cap S^{1\times n}$. Then $x = uA$ for some $u \in T^{1\times m}$. We want to show that $x \in \mathrm{Ker}^2_S \mathrm{Row}_S(A)$, so let $v, v' \in S^{n\times 1}$ and suppose that $Av = Av'$. Then $xv = uAv = uAv' = xv'$, and thus $(v, v') \in \mathrm{Ker}_S(x)$. Therefore $x \in \mathrm{Ker}^2_S \mathrm{Row}_S(A)$ because we have just shown that $\mathrm{Ker}_S \mathrm{Row}_S(A) \subseteq \mathrm{Ker}_S(x)$. Hence $x \in \mathrm{Row}_S(A)$ by Proposition 9.10.

(ii). First suppose that $\mathrm{Row}_S(A) \subseteq \mathrm{Row}_S(B)$. Then, in particular, $\mathrm{Row}_S(B)$ contains each row of A. Therefore $\mathrm{Row}_T(B)$ contains each row of A because $S \subseteq T$, and thus $\mathrm{Row}_T(A) \subseteq \mathrm{Row}_T(B)$ because, by Proposition 6.5 dual, $\mathrm{Row}_T(A)$ is the smallest left T-submodule of $T^{1\times n}$ that contains each row of A. Conversely, if $\mathrm{Row}_T(A) \subseteq \mathrm{Row}_T(B)$ then $\mathrm{Row}_T(A) \cap S^{1\times n} \subseteq \mathrm{Row}_T(B) \cap S^{1\times n}$, and so we conclude that $\mathrm{Row}_S(A) \subseteq \mathrm{Row}_S(B)$ by (i). □

Lemma 12.1 (ii) tells us that if S is a right exact subsemiring of a semiring T then the function given by $\mathrm{Row}_S(A) \mapsto \mathrm{Row}_T(A)$ is an order embedding from the poset of row spaces in $S^{1\times n}$ to the poset of row spaces in $T^{1\times n}$ (see page 93). This function is well-defined and injective because we have $\mathrm{Row}_S(A) = \mathrm{Row}_S(B)$ if and only if $\mathrm{Row}_T(A) = \mathrm{Row}_T(B)$ for all $A \in S^{m\times n}$ and all $B \in S^{p\times n}$, and the fact that it is an order embedding means that the inclusion structure of row spaces over T must be at least as complicated as the inclusion structure of row spaces over S. In particular, T cannot have a simpler poset of finitely generated left ideals than S.

Theorem 12.2 *Let T be a semiring and let S be a subsemiring of T. If S is right exact then there is an order embedding from the poset of finitely generated left ideals of S to the poset of finitely generated left ideals of T.*

Proof Apply Lemma 12.1 (ii) in the case $n = 1$, as described above. □

One consequence of Theorem 12.2 is that any exact commutative non-trivial ring that is a subring of a field must itself be a field. This is because a field only has two finitely generated ideals: the zero ideal and the whole field. An exact non-trivial subring of a field could not have more than two finitely generated ideals as there would be no order embedding, and yet it must have at least two finitely generated ideals because otherwise it would then be the trivial ring. Such a ring must therefore have precisely two finitely generated ideals, so must be a field. Every integral domain is a subring of a field (see Cohn [21, Theorem 6.2.3]), and thus by the above observation an exact integral domain must be a field. We also outlined a direct proof of this fact just after Example 9.3.

Now—slightly tangentially—recall that our original definition of exactness was that a semiring is right exact if each right S-linear function from the column space of a matrix to S has an extension (see Definition 9.2). In particular, if S is right exact then every linear functional on a finitely generated right ideal of S has an extension. Does this property characterise right exactness? That is, given a matrix $A \in S^{m\times n}$ and a linear function $f\colon \mathrm{Col}(A) \to S$, is it possible to produce an extension of f only using extensions of linear functionals on finitely generated right ideals of S? If so, it would be sufficient to define exactness in terms of linear functionals on finitely generated ideals.

The corresponding question for self-injectivity of rings has a positive answer: by the work of Baer [5], a ring R is right self-injective if and only if each linear functional on a right ideal of R has an extension (see Rotman [73, Theorem 3.30]). This means that (when we are concerned with injectivity, at least) linear functionals on right ideals of R capture the complexity of linear functionals on right R-modules. Whether this works for exactness—also known as FP-injectivity—of rings is an open question, however. The property of all linear functionals on finitely generated ideals of a ring R having an extension is known as *F-injectivity* of R, and it is unknown whether every F-injective ring is FP-injective (see Nicholson and Yousif [70, Question 10]).

It is much easier to resolve the above question for semirings. In Example 20.6 we will construct a non-exact semiring with the property that each linear functional

on a finitely generated ideal has an extension. This means that, in general, it is not possible to transfer partial exactness up from finitely generated ideals of a semiring to row and column spaces of matrices. It turns out that Bear's test for self-injectivity does not work when stated for semirings either. For instance, the semiring in Example 20.6 is finite, and as such each ideal has the linear functional extension property, but (by Proposition 9.5) this semiring cannot be self-injective as it is not even exact.

If S is an exact subsemiring of a semiring T then Lemma 12.1 restricts parts of the structure of matrices over T, but it does not tell us anything about the parts that do not intersect with S. Away from S, the behaviour of matrices could be completely different, so it seems unlikely that we could formulate a useful characterisation of when exactness can be transferred up from S to T. The opposite direction in which we could consider transferring exactness is, of course, down from a semiring to a subsemiring. We are most interested in when it is possible to transfer exactness down from a full matrix semiring to a subsemiring, so we begin by considering what happens for matrix semirings comprising all the matrices of a given shape. For instance, we might wonder if semirings of upper triangular matrices can ever be exact. To make this notion precise, we introduce the following definitions.

Let S be a standard semiring. A relation $E \subseteq \{(i,j) : 1 \leq i,j \leq n\}$ will be called an *n-shape* if it is reflexive and transitive, and given such a relation we define

$$\mathrm{M}_E(S) = \{A \in \mathrm{M}_n(S) : A_{ij} = 0 \text{ for all } (i,j) \notin E\}. \tag{12.1}$$

An n-shape E records the matrix entries that are permitted to be non-zero, so $\mathrm{M}_E(S)$ should be interpreted as comprising all the matrices of shape E. For example, if we take $E = \{(i,i) : 1 \leq i \leq n\}$ then $\mathrm{M}_E(S)$ is the semiring of diagonal $n \times n$ matrices with entries in S, and if we take $E = \{(i,j) : 1 \leq i \leq j \leq n\}$ then $\mathrm{M}_E(S)$ is the semiring of upper triangular $n \times n$ matrices with entries in S. Notice that if we take $E = \{(i,j) : 1 \leq i,j \leq n\}$ then we simply recover $\mathrm{M}_n(S)$ because (12.1) does not force any matrix entries to be zero.

The definition of an n-shape ensures that $\mathrm{M}_E(S)$ is a subsemiring of $\mathrm{M}_n(S)$ for each n-shape E. Specifically, the assumption that E is transitive means that $\mathrm{M}_E(S)$ is closed under multiplication, and the assumption that E is reflexive means that $\mathrm{M}_E(S)$ contains the identity matrix. (It is clear from (12.1) that $\mathrm{M}_E(S)$ is always going to be an additive submonoid of $\mathrm{M}_n(S)$, so we need not worry about

addition or zero.) We are now in a position to characterise the shapes for which the corresponding matrix semiring inherits exactness from S.

Theorem 12.3 *Let S be a right exact standard semiring and let E be an n-shape. Then $\mathrm{M}_E(S)$ is right exact if and only if E is an equivalence relation.*

Proof ($\Rightarrow$). Suppose that $\mathrm{M}_E(S)$ is right exact. By definition, an n-shape is reflexive and transitive, so it is sufficient to show that E is symmetric. Also notice that $(i,j) \in E$ if and only if $\delta_{ij} \in \mathrm{M}_E(S)$, where $\delta_{ij} \in \mathrm{M}_n(S)$ denotes the matrix with i-jth entry 1 and all other entries 0. It is therefore enough to show that $\delta_{ij} \in \mathrm{M}_E(S)$ implies that $\delta_{ji} \in \mathrm{M}_E(S)$ for all $1 \leq i, j \leq n$.

Let $1 \leq i, j \leq n$ and suppose that $\delta_{ij} \in \mathrm{M}_E(S)$. Since E is reflexive we have $\delta_{ji}\delta_{ij} = \delta_{jj} \in \mathrm{M}_E(S)$, and thus $\delta_{jj} \in \mathrm{M}_n(S)\delta_{ij} \cap \mathrm{M}_E(S)$. Lemma 12.1 (i) then gives $\delta_{jj} \in \mathrm{M}_E(S)\delta_{ij}$, because $\mathrm{M}_E(S)$ is right exact with $\delta_{ij} \in \mathrm{M}_E(S)$, and as such there is some $A \in \mathrm{M}_E(S)$ satisfying $\delta_{jj} = A\delta_{ij}$. Therefore

$$\delta_{ji} = \delta_{jj}\delta_{ji} = A\delta_{ij}\delta_{ji} = A\delta_{ii} \in \mathrm{M}_E(S) \tag{12.2}$$

because $A, \delta_{ii} \in \mathrm{M}_E(S)$. Hence E is symmetric.

($\Leftarrow$). Suppose that E is symmetric. Then E has $1 \leq m \leq n$ equivalence classes $E_1, \ldots, E_m$ with cardinalities $n_1, \ldots, n_m$ (respectively) satisfying $n_1 + \cdots + n_m = n$. Each class E_k has a corresponding n-shape $E_k \times E_k$, and since the E_k are pairwise disjoint we can identify $\mathrm{M}_E(S)$ with the direct product

$$\mathrm{M}_{E_1 \times E_1}(S) \times \cdots \times \mathrm{M}_{E_m \times E_m}(S). \tag{12.3}$$

Using a bijection between E_k and $\{1, \ldots, n_k\}$, we can then identify $\mathrm{M}_{E_k \times E_k}(S)$ with $\mathrm{M}_{n_k}(S)$, and thus we can identify $\mathrm{M}_E(S)$ with the direct product

$$\mathrm{M}_{n_1}(S) \times \cdots \times \mathrm{M}_{n_m}(S). \tag{12.4}$$

By assumption S is right exact, so by Proposition 11.1 each $\mathrm{M}_{n_k}(S)$ is right exact. Hence $\mathrm{M}_E(S)$ is right exact by Proposition 11.2. □

Theorem 12.3 tells us everything we really want to know about when the semiring of matrices of a given shape over an exact standard semiring S is exact. Essentially, the shape must be symmetric in order for the corresponding matrix semiring to be

exact. For example then, semirings of upper triangular matrices will never be exact, but the semiring comprising all the matrices of shape

$$\begin{bmatrix} a & 0 & b \\ 0 & c & 0 \\ d & 0 & e \end{bmatrix}, \tag{12.5}$$

for $a, b, c, d, e \in S$, will be exact because it can be identified with $\mathrm{M}_2(S) \times S$.

In general, a full matrix semiring will have more subsemirings that just those comprising all the matrices of a given shape. For instance, if S is a semiring and G is a finite group of order n then Lemma 7.2 tells us that the group semiring SG can be viewed as a subsemiring of $\mathrm{M}_n(S)$. It would therefore be helpful if we could find sufficient conditions for an arbitrary matrix semiring to inherit exactness from $\mathrm{M}_n(S)$. This problem of transferring exactness down from a semiring to a subsemiring appears more amenable to a general result than the problem of transferring exactness up from a subsemiring, as there are no unknown elements to account for. Indeed, the following result shows that if a subsemiring is sufficiently similar to its containing semiring then downwards transfer of exactness is possible.

Theorem 12.4 *Let T be a right exact semiring and let S be a right retract of T. If there is an injective left S-linear function $g\colon T \to S^{1\times q}$ for some q then S is right exact.*

Proof Let $A \in S^{m\times n}$ and let $y \in \mathrm{Ker}^2_S \mathrm{Row}_S(A)$. Then $\mathrm{Ker}_S \mathrm{Row}_S(A) \subseteq \mathrm{Ker}_S(y)$. We begin by showing that $y \in \mathrm{Ker}^2_T \mathrm{Row}_T(A)$. To do this, let $v, v' \in T^{n\times 1}$ and suppose that $Av = Av'$. By applying g entrywise, we can extend it to a function $g\colon T^{n\times 1} \to S^{n\times q}$, and the assumption that g is left S-linear then means that we have $g(Av) = A(gv)$ and $g(Av') = A(gv')$. Therefore $A(gv)w = A(gv')w$ for all $w \in S^{q\times 1}$ because $Av = Av'$, and as such $((gv)w, (gv')w) \in \mathrm{Ker}_S \mathrm{Row}_S(A) \subseteq \mathrm{Ker}_S(y)$ for all $w \in S^{q\times 1}$. This means that $y(gv)w = y(gv')w$ for all $w \in S^{q\times 1}$, so by Proposition 6.3 we have $y(gv) = y(gv')$. Left S-linearity of g then gives $g(yv) = g(yv')$, and thus $yv = yv'$ because g is injective. That is, $(v, v') \in \mathrm{Ker}_T(y)$. We have therefore shown that $\mathrm{Ker}_T \mathrm{Row}_T(A) \subseteq \mathrm{Ker}_T(y)$, and as such $y \in \mathrm{Ker}^2_T \mathrm{Row}_T(A)$.

We can now use right exactness of T. By Proposition 9.10 we have $y \in \mathrm{Row}_T(A)$, which means that $y = uA$ for some $u \in T^{1\times m}$. Since S is a right retract of T, there is a right S-linear function $f\colon T \to S$ that fixes S pointwise (see Definition 5.7).

Note that, as above, f can be extended to a function $f\colon T^{1\times m} \to S^{1\times m}$ and to a function $f\colon T^{1\times n} \to S^{1\times n}$. The assumption that f is S-linear and fixes S pointwise then gives $y = fy = f(uA) = (fu)A$, and as such $y \in \mathrm{Row}_S(A)$. We have therefore shown that $\mathrm{Ker}_S^2 \mathrm{Row}_S(A) \subseteq \mathrm{Row}_S(A)$, and so in fact $\mathrm{Ker}_S^2 \mathrm{Row}_S(A) = \mathrm{Row}_S(A)$. Hence S is right exact by Proposition 9.10. □

Corollary 12.5 *If S is a right exact semiring then the group semiring SG is right exact for every finite group G.*

Proof Suppose that G has order n. By Theorem 7.3 (iii), SG is a right retract of $\mathrm{M}_n(S)$, and by Theorem 7.3 (ii) dual, there is (in particular) an injective left SG-linear function $\mathrm{M}_n(S) \to (SG)^{1\times n}$. Hence SG is right exact by Proposition 11.1 and Theorem 12.4. □

Shitov [76, Theorem 3.5] has shown that if R is a ring and G is a group then the group ring RG is exact if and only if R is exact and G is locally finite. His proof that exactness of RG implies exactness of R essentially uses Theorem 12.4 to transfer exactness down from RG to R. Specifically, R is a right retract of RG via the function that selects any diagonal entry (i.e., the constant term) of an element of RG, and the required injective left R-linear function $RG \to R^{1\times n}$ is given by selecting the first row of an element of RG.

Linear algebra over rings

13 Orthogonal complements and exact annihilators

The kernel of a set of row vectors over a semiring records the pairs of column vectors that give the same result when multiplied with each row vector in the set (see Definition 8.1). Over a ring, however, we do not need to keep track of pairs of column vectors because we can rewrite an equality of the form $xv = xv'$ as $x(v - v') = 0$, and this means that it suffices to simply record which single column vectors give the zero vector when multiplied with each row vector. Consequently, the kernel becomes much easier to understand, as it is sufficient to only consider one of its equivalence classes. This representative class is called the 'orthogonal complement' of the set of row vectors.

In this section we recall some basic properties of orthogonal complements and we show that exactness of a ring is equivalent to a familiar double orthogonal complement condition. Orthogonal complements are closely linked with annihilators, so we also introduce a notion of 'exact annihilator' and the corresponding property of a ring being an 'exact annihilator ring'. Being an exact annihilator ring is slightly stronger than being an exact ring, and this means that, in addition to exactness, we can gain valuable information about the structure of orthogonal complements by constructing exact annihilators. We will illustrate this for various rings in section 14.

Definition 13.1 Let R be a ring and let $X \subseteq R^{1\times n}$. The *orthogonal complement* of X is the set

$$X^{\perp} = \big\{v \in R^{n\times 1} : xv = 0 \text{ for all } x \in X\big\}. \tag{13.1}$$

Comparing this definition with Definition 8.1, we see that $X^{\perp}$ is precisely the equivalence class $[0]_{\mathrm{Ker}(X)}$. As mentioned above, in the context of rings this is the only class we need to consider because if $v, v' \in R^{n\times 1}$ then $(v, v') \in \mathrm{Ker}(X)$ if and only if $v - v' \in X^{\perp}$. Put another way, since $X^{\perp}$ is a right R-submodule of $R^{n\times 1}$,

it is (in particular) an additive subgroup of $R^{n\times 1}$, and the equivalence classes of $\operatorname{Ker}(X)$ are just the cosets of $X^\perp$. This means that quotienting $R^{n\times 1}$ by the right R-submodule $X^\perp$ gives the same result as quotienting it by the right R-congruence $\operatorname{Ker}(X)$ would, and so Proposition 8.6 (ii) immediately gives Proposition 13.2 (i), below.

Part (ii) of Proposition 13.2 follows from the fact that, as with kernels, taking orthogonal complements constitutes a Galois connection between $R^{1\times n}$ and $R^{n\times 1}$. The only difference in the case of orthogonal complements is that both functions of the Galois connection take a set of vectors to a set of vectors, rather than one of them taking a relation on vectors to a set of vectors instead (see Proposition 8.4). That is, a set of column vectors also has an orthogonal complement—defined dually to (13.1)—and together the two functions between $\operatorname{Pow}\bigl(R^{1\times n}\bigr)$ and $\operatorname{Pow}\bigl(R^{n\times 1}\bigr)$ form a Galois connection.

Proposition 13.2 *Let R be a ring and let $A \in R^{m\times n}$. Then*

(i) $R^{n\times 1}/\operatorname{Row}(A)^\perp \cong \operatorname{Col}(A)$ *as right R-modules; and*

(ii) $\operatorname{Row}(A)^{\perp\perp\perp} = \operatorname{Row}(A) \subseteq \operatorname{Row}(A)^{\perp\perp}$.

Proof See above. □

Since it is simpler and sufficient (when working with rings) to study orthogonal complements instead of kernels, we would like to turn Proposition 9.10 into a characterisation of exactness that is phrased purely in terms of orthogonal complements. In view of Proposition 13.2 (ii), and the fact that Proposition 9.10 involves a double kernel, the following result is the appropriate such characterisation.

Proposition 13.3 *A ring R is right exact if and only if* $\operatorname{Row}(A)^{\perp\perp} = \operatorname{Row}(A)$ *for all $A \in R^{m\times n}$.*

Proof By Proposition 9.10 it is enough to show that $\operatorname{Row}(A)^{\perp\perp} = \operatorname{Ker}^2 \operatorname{Row}(A)$ for all $A \in R^{m\times n}$, but in fact this holds more generally because if $X \subseteq R^{1\times n}$ then we have

$$\begin{aligned} X^{\perp\perp} &= \bigl\{y \in R^{1\times n} : yv = 0 \text{ for all } v \in X^\perp\bigr\} \\ &= \bigl\{y \in R^{1\times n} : X^\perp \subseteq \{y\}^\perp\bigr\} \\ &= \bigl\{y \in R^{1\times n} : \operatorname{Ker}(X) \subseteq \operatorname{Ker}(y)\bigr\} \\ &= \operatorname{Ker}^2(X) \end{aligned} \tag{13.2}$$

by Definition 13.1 dual and Definition 8.3. □

We will not actually use Proposition 13.3 to show that any specific rings are exact, however. Instead, the rings considered in section 14 will be shown to be exact using a stronger condition on annihilators of matrices (see Proposition 13.7, below). A *right annihilator* of a matrix $A \in R^{m\times n}$ is simply a matrix $B \in R^{n\times q}$ satisfying $AB = 0$, but for our purposes the following characterisation will be more useful.

Proposition 13.4 *Let R be a ring, let $A \in R^{m\times n}$ and let $B \in R^{n\times q}$. Then B is a right annihilator of A if and only if* $\mathrm{Row}(A) \subseteq \mathrm{Col}(B)^{\perp}$.

Proof Following Proposition 8.6 (i) dual, the definition of $\mathrm{Col}(B)^{\perp}$ can be simplified to

$$\mathrm{Col}(B)^{\perp} = \big\{x \in R^{1\times n} : xB = 0\big\}. \tag{13.3}$$

It is apparent from this description that $\mathrm{Col}(B)^{\perp}$ is what would sometimes be called the 'left null space' of B.

($\Rightarrow$). Suppose that $AB = 0$ and let $x \in \mathrm{Row}(A)$. Then $x = uA$ for some $u \in R^{1\times m}$. Therefore $xB = uAB = 0$, and as such $x \in \mathrm{Col}(B)^{\perp}$ by (13.3).

($\Leftarrow$). Suppose that $\mathrm{Row}(A) \subseteq \mathrm{Col}(B)^{\perp}$. Then $uAB = 0$ for all $u \in R^{1\times m}$, by (13.3) again, and thus $AB = 0$. □

Each matrix $A \in R^{m\times n}$ has a collection of trivial right annihilators—namely the zero matrices of all appropriate sizes—and if A happens to be invertible then it is obvious that these are the only right annihilators of A. Even if A is not invertible, it is still possible for the zero matrices to be the only annihilators of A (e.g., consider the non-invertible "matrix" $2 \in \mathbf{Z}$, as in Example 9.3), but if R is exact then there is a very satisfying relationship between inverses and annihilators. The following generalisation of a result of Wilding et al. [86, Proposition 4.2] makes this relationship precise.

Proposition 13.5 *Let R be a right exact ring and let $A \in R^{m\times n}$. Then A has a left inverse if and only if A has no non-zero right annihilator.*

Proof ($\Rightarrow$). Suppose that A has a left inverse $P \in R^{n\times m}$ and let $B \in R^{n\times q}$ be a right annihilator of A. Then we have $B = I_n B = PAB = P0 = 0$, where I_n denotes the $n \times n$ identity matrix, and consequently A has no non-zero right annihilator.

($\Leftarrow$). Suppose that A has no non-zero right annihilator. Then $\mathrm{Row}(A)^{\perp} = \{0\}$ because each $v \in \mathrm{Row}(A)^{\perp}$ is a right annihilator of A, and thus $\mathrm{Row}(A)^{\perp\perp} = R^{1\times n}$.

Therefore $\mathrm{Row}(A) = \mathrm{Row}(A)^{\perp\perp} = R^{1\times n} = \mathrm{Row}(I_n)$ by Proposition 13.3, since R is right exact, and as such there is some $P \in R^{n\times m}$ with $PA = I_n$ by Proposition 6.6 dual. □

Now recall that our general strategy for showing that an arbitrary semiring is right exact is to find a relation F_A satisfying $\mathrm{Row}(A) = \mathrm{Ker}(F_A)$ for each matrix A with entries in S (see Lemma 9.12). In the context of rings, a more fruitful approach would be to try to find a matrix $B_A \in R^{n\times q}$ with $\mathrm{Row}(A) = \mathrm{Col}(B_A)^{\perp}$ for each $A \in R^{m\times n}$, and it is clear from Proposition 13.3 that this would mean looking for a special kind of right annihilator of A. By Proposition 13.2 (ii) dual, such a matrix B_A would give

$$\mathrm{Row}(A)^{\perp\perp} = \mathrm{Col}(B_A)^{\perp\perp\perp} = \mathrm{Col}(B_A)^{\perp} = \mathrm{Row}(A), \tag{13.4}$$

and thus if each $A \in R^{m\times n}$ has one of these special right annihilators then R is right exact by Proposition 13.3. This observation motivates the following definitions and proves Proposition 13.7, below.

Definition 13.6 Let R be a ring and let $A \in R^{m\times n}$. A *right exact annihilator* of A is a matrix $B \in R^{n\times q}$ satisfying $\mathrm{Row}(A) = \mathrm{Col}(B)^{\perp}$.

If each $A \in R^{m\times n}$ has a right exact annihilator then we will call R a *right exact annihilator ring*, and, dually, if each $A \in R^{m\times n}$ has a left exact annihilator then we will call R a *left exact annihilator ring*. Note that a *left exact annihilator* of A is a matrix $B \in R^{p\times m}$ satisfying $\mathrm{Col}(A) = \mathrm{Row}(B)^{\perp}$. A ring which is both a left and a right exact annihilator ring will be called an *exact annihilator ring*.

Proposition 13.7 *If R is a right exact annihilator ring then R is right exact.*

Proof Apply Proposition 13.2 (ii) dual and Proposition 13.3 to Definition 13.6, as described above. □

The definition of an exact annihilator illustrates why we chose to use the term 'exact' in the first place: if $B \in R^{n\times q}$ is a right exact annihilator of $A \in R^{m\times n}$ then we have an exact sequence

$$R^{1\times m} \xrightarrow{\quad A \quad} R^{1\times n} \xrightarrow{\quad B \quad} R^{1\times q} \tag{13.5}$$

of left R-modules. Specifically, $\mathrm{Row}(A)$ is the image of the function $u \mapsto uA$ and, by (13.3), $\mathrm{Col}(B)^\perp$ is the kernel of the function $x \mapsto xB$ (in the standard R-module sense), so if B is a right exact annihilator of A then the image of the first function equals the kernel of the second function because $\mathrm{Row}(A) = \mathrm{Col}(B)^\perp$. This makes (13.5) an exact sequence.

The existence of a right exact annihilator of a matrix $A \in R^{m\times n}$ gives us some information about $\mathrm{Row}(A)^\perp$, as we have $\mathrm{Row}(A)^\perp = \mathrm{Col}(B)^{\perp\perp}$ whenever B is a right exact annihilator of A, but in general we will not really know anything about the structure of $\mathrm{Col}(B)^{\perp\perp}$ unless R is left exact. Dually, if we start with a left exact annihilator ring then we might not be able to fully understand $\mathrm{Col}(A)^\perp$ unless R is right exact. We will therefore be in the best position to describe $\mathrm{Row}(A)^\perp$ and $\mathrm{Col}(A)^\perp$ if we start with an exact annihilator ring.

Proposition 13.8 *Let R be an exact annihilator ring and let $A \in R^{m\times n}$. Then there is some $B \in R^{n\times q}$ satisfying* $\mathrm{Row}(A) = \mathrm{Col}(B)^\perp$ *and* $\mathrm{Row}(A)^\perp = \mathrm{Col}(B)$.

Proof Since R is a right exact annihilator ring there is some $B \in R^{n\times q}$ satisfying $\mathrm{Row}(A) = \mathrm{Col}(B)^\perp$. Therefore $\mathrm{Row}(A)^\perp = \mathrm{Col}(B)^{\perp\perp}$, as above. Now since R is also a left exact annihilator ring it is, in particular, left exact by Proposition 13.7 dual, and thus $\mathrm{Row}(A)^\perp = \mathrm{Col}(B)$ by Proposition 13.3 dual. □

Proposition 13.8 tells us that if our aim is to understand orthogonal complements over a ring R (which it is) then we should try to show that R is an exact annihilator ring. Moreover, it also tells us that it is the exact annihilators themselves that are key to gaining this understanding. The focus of the next section will therefore be on how to actually construct exact annihilators over specific rings.

14 Commutative elementary divisor rings

The purpose of this section is to give examples of rings that behave very much like fields—at least from the point of view of linear algebra. The results in section 13 suggest that exact annihilator rings are good candidates because, for one thing, they are exact, but also because they admit a promising description of orthogonal complements in terms of exact annihilators. To simplify the process of finding exact annihilators, our basic assumption will be that we have an elementary divisor ring. This will allow us to reduce the problem of constructing exact annihilators

of arbitrary matrices to the much less complicated problem of constructing exact annihilators of elements of the ring.

The examples we consider in this section (homomorphic images of principal ideal domains and, later, Boolean rings) are commutative, so we will usually assume that we are working with a commutative elementary divisor ring. As we show in Theorem 14.2, the row and column spaces of a matrix with entries in such a ring are actually isomorphic, and this leads to a very nice relationship between the row space of a matrix and its orthogonal complement.

Definition 14.1 A ring R is an *elementary divisor ring* if for each $A \in R^{m\times n}$ there are invertible matrices $P \in R^{m\times m}$ and $Q \in R^{n\times n}$ with PAQ diagonal.

Examples of elementary divisor rings include the integers (see below) and the ring of real analytic functions (see Brewer et al. [12, pages 96–100]). The above definition of an elementary divisor ring was introduced by Kaplansky [49], but the idea goes back at least as far as Smith [81] who introduced a procedure to find invertible matrices $P \in \mathbf{Z}^{m\times m}$ and $Q \in \mathbf{Z}^{n\times n}$ with PAQ diagonal for any given $A \in \mathbf{Z}^{m\times n}$.[1] This procedure is not special to the integers though; it also works for matrices with entries in any *principal ideal domain* (an integral domain in which each ideal is generated by a single element). That is, every principal ideal domain is an elementary divisor ring (see Cohn [21, Theorem 10.5.4]). By definition, a principal ideal domain is also commutative, and as such the following result applies to every principal ideal domain.

Theorem 14.2 *If R is a commutative elementary divisor ring then*

(i) *the identity function on R is a conjugation; and*

(ii) $\mathrm{Row}(A) \cong \mathrm{Col}(A)$ *as R-modules for all $A \in R^{m\times n}$.*

Proof (i). Let $A \in R^{n\times n}$. Then since R is an elementary divisor ring there are invertible matrices $P \in R^{m\times m}$ and $Q \in R^{n\times n}$ with PAQ square (because A is square) diagonal. In particular then, $(PAQ)^{\mathrm{T}} = PAQ$. Commutativity of R implies that the identity function on R is a standard involution, which in turn means that

$$A\big(P^{-1}Q^{\mathrm{T}}\big)^{\mathrm{T}}P^{\mathrm{T}}Q^{-1} = AQP^{-\mathrm{T}}P^{\mathrm{T}}Q^{-1} = AQQ^{-1} = A. \tag{14.1}$$

[1] Kaplansky [49] originally also placed a divisibility condition on the diagonal entries of PAQ, but at some point it appears to have become acceptable to drop this condition.

Therefore the matrices $M = P^{-1}Q^{\mathrm{T}}$ and $N = P^{\mathrm{T}}Q^{-1}$ satisfy one of the requirements in the hypothesis of Lemma 10.3. For the other requirement to be satisfied, we need to show that $MA^{\mathrm{T}} = AM^{\mathrm{T}}$, but this holds because we have

$$P^{-1}Q^{\mathrm{T}}A^{\mathrm{T}} = P^{-1}\big(Q^{\mathrm{T}}A^{\mathrm{T}}P^{\mathrm{T}}\big)P^{-\mathrm{T}} = P^{-1}(PAQ)P^{-\mathrm{T}} = AQP^{-\mathrm{T}}. \tag{14.2}$$

Hence the identity function on R is a conjugation by Lemma 10.3.

(ii). As noted above, commutativity of R implies that the identity function on R is a standard involution. By (i) it is also a conjugation, so by Proposition 10.2 there is an R-module isomorphism $\operatorname{Row}(A) \to \operatorname{Col}(A)$ for each $A \in R^{m\times n}$. □

In particular, Theorem 14.2 (ii) tells us that the row and column spaces of any given integer matrix are isomorphic. This fact is interesting, but not especially useful because the integers are our standard example of a non-exact ring (see Example 9.3). The real benefit of having an isomorphism between the row and column spaces of a matrix becomes clear when we consider exact annihilators, and, of course, $\mathbf{Z}$ is not an exact annihilator ring.

Theorem 14.3 *Let R be a commutative elementary divisor ring. If R is an exact annihilator ring then $\operatorname{Row}(A)^{\perp} \cong R^{1\times n}/\operatorname{Row}(A)$ as R-modules for all $A \in R^{m\times n}$.*

Proof Suppose that R is an exact annihilator ring and let $A \in R^{m\times n}$. Then by Proposition 13.8 there is some $B \in R^{n\times q}$ with $\operatorname{Row}(A) = \operatorname{Col}(B)^{\perp}$ and $\operatorname{Row}(A)^{\perp} = \operatorname{Col}(B)$. Therefore $\operatorname{Row}(A)^{\perp} \cong \operatorname{Row}(B) \cong R^{1\times n}/\operatorname{Col}(B)^{\perp}$ by Theorem 14.2 (ii) and Proposition 13.2 (i) dual. Hence $\operatorname{Row}(A)^{\perp} \cong R^{1\times n}/\operatorname{Row}(A)$ as R-modules because $\operatorname{Row}(A) = \operatorname{Col}(B)^{\perp}$. □

This property of row spaces is well-known in the case R is a field, as it follows from standard results concerning rank and nullity (see Roman [72, Theorems 1.16 and 2.8]). For example, if $R = \mathbf{R}$ and $A = \left[\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{smallmatrix}\right]$ then $\operatorname{Row}(A)$ is a plane passing through the origin in three-dimensional space and the cosets of this plane lie along a perpendicular line passing through the origin. The orthogonal complement of $\operatorname{Row}(A)$ is also a line in three-dimensional space because it has basis $\left[\begin{smallmatrix} 0 & 0 & 1 \end{smallmatrix}\right]^{\mathrm{T}}$, and thus $\operatorname{Row}(A)^{\perp} \cong \mathbf{R}^{1\times 3}/\operatorname{Row}(A)$ as vector spaces.

As mentioned above, our main reason for considering elementary divisor rings is to make it as easy as possible to construct exact annihilators of matrices. The proof of the following result demonstrates how to use the existence of exact annihilators

of elements of an elementary divisor ring to construct an exact annihilator of an arbitrary matrix. We will illustrate this procedure in Examples 14.6 and 14.8.

Lemma 14.4 *Let R be an elementary divisor ring. If each element of R has a right exact annihilator then R is a right exact annihilator ring.*

Proof Suppose that each element of R has a right exact annihilator. Then by Definition 13.6, for each $a \in R$ there is some row vector u with $Ra = \mathrm{Col}(u)^{\perp}$.

We begin by showing that each diagonal matrix $D \in R^{m\times n}$ has a right exact annihilator. To do this, we need to find some matrix C with $\mathrm{Row}(D) = \mathrm{Col}(C)^{\perp}$, and so since adding or removing zero rows does not change $\mathrm{Row}(D)$ we may assume that D is square diagonal. That is, we may assume that D has diagonal entries $a_1, \ldots, a_n \in R$. We then define $C \in R^{n\times q}$ by

$$C = \begin{bmatrix} u_1 & 0 & \ldots & 0 \\ 0 & u_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & u_n \end{bmatrix}, \tag{14.3}$$

where each $u_i \in R^{1\times q_i}$ satisfies $Ra_i = \mathrm{Col}(u_i)^{\perp}$, and where $q = q_1 + \cdots + q_n$. By construction, this matrix satisfies

$$\mathrm{Col}(C)^{\perp} = \big\{x \in R^{1\times n} : x_{1i} \in \mathrm{Col}(u_i)^{\perp} \text{ for all } 1 \leq i \leq n\big\}, \tag{14.4}$$

and since D is diagonal we also have

$$\mathrm{Row}(D) = \big\{x \in R^{1\times n} : x_{1i} \in Ra_i \text{ for all } 1 \leq i \leq n\big\}. \tag{14.5}$$

Therefore $\mathrm{Row}(D) = \mathrm{Col}(C)^{\perp}$ because $Ra_i = \mathrm{Col}(u_i)^{\perp}$ for each $1 \leq i \leq n$, and as such C is a right exact annihilator of D.

Now let $A \in R^{m\times n}$ be an arbitrary matrix with entries in R. Since R is an elementary divisor ring there are invertible matrices $P \in R^{m\times m}$ and $Q \in R^{n\times n}$ with PAQ diagonal, so by the above argument there is some $C \in R^{n\times q}$ satisfying $\mathrm{Row}(PAQ) = \mathrm{Col}(C)^{\perp}$. The fact that P and Q are invertible means that

$$\mathrm{Row}(A) = \big\{x \in R^{1\times n} : xQ \in \mathrm{Row}(PAQ)\big\}, \tag{14.6}$$

and thus $x \in \mathrm{Row}(A)$ if and only if $xQ \in \mathrm{Col}(C)^{\perp}$. Therefore $\mathrm{Row}(A) = \mathrm{Col}(QC)^{\perp}$, as it is obvious that $xQ \in \mathrm{Col}(C)^{\perp}$ if and only if $x \in \mathrm{Col}(QC)^{\perp}$, and as such QC is a right exact annihilator of A. Hence R is a right exact annihilator ring. □

Now recall that while every principal ideal domain is a commutative elementary divisor ring (see page 82), not every principal ideal domain is an exact annihilator ring (e.g., the ring $\mathbf{Z}$ is not even exact; see Example 9.3). As a result, Theorem 14.3 is not applicable to all principal ideal domains. However, it turns out that all proper quotients of principal ideal domains are exact annihilator rings (in addition to being commutative elementary divisor rings), so in particular Theorem 14.3 is applicable to the ring $\mathbf{Z}/m\mathbf{Z}$ for each $m > 0$ because the following result is.

Theorem 14.5 *If R is a proper homomorphic image of a principal ideal domain then*

(i) *R is a commutative elementary divisor ring; and*

(ii) *R is an exact annihilator ring.*

Proof Since R is a proper homomorphic image of a principal ideal domain, there is a principal ideal domain R' and a surjective homomorphism $f\colon R' \to R$ with $\mathrm{Ker}(f) \neq \{0\}$.

(i). As Henriksen [38, Proof of Theorem 3] notes, any homomorphic image of a (commutative) elementary divisor ring is again a (commutative) elementary divisor ring. Therefore R is a commutative elementary divisor ring because R' is a principal ideal domain.

(ii). Since R is commutative it is sufficient to show that R is a right exact annihilator ring, because if B is a right exact annihilator of $A \in R^{m\times n}$ then B^{T} is a left exact annihilator of A^{T}. By (i) R is an elementary divisor ring, so by Lemma 14.4 it is actually enough to show that each $a \in R$ has a right exact annihilator, but now since $a^{\mathrm{T}} = a$ we do not need to worry about distinguishing between left and right exact annihilators of a.

Let $a \in R$. Then since f is surjective there is some $a' \in R'$ with $fa' = a$. The kernel of f is an ideal of R', so since R' is a principal ideal domain there is some $d' \in R'$ with $\mathrm{Ker}(f) = d'R'$. Moreover, we have $d' \neq 0$ because by assumption $\mathrm{Ker}(f) \neq \{0\}$. Now take $r \in R'$ to generate the ideal $a'R' + d'R'$, then we can write

$a' = rs$ and $d' = rb'$ for some $s, b' \in R'$. Note that $b' \neq 0$ because $d' \neq 0$. If we take $b = fb'$ then we have

$$ab = (fa')(fb') = f(a'b') = f(rsb') = f(rb's) = f(d's) = 0 \tag{14.7}$$

because f is a homomorphism and because $d's \in d'R' = \mathrm{Ker}(f)$. Therefore b is an annihilator of a.

To show that b is an exact annihilator of a, let $c \in R$ and suppose that $bc = 0$. We then want $c \in aR$. Writing $c = fc'$ for some $c' \in R'$, we have $b'c' \in \mathrm{Ker}(f) = d'R'$ because $f(b'c') = (fb')(fc') = bc = 0$, and thus $b'c' = d't = rb't$ for some $t \in R'$. Since R' is an integral domain (and since $b' \neq 0$) this implies that $c' = rt$. Therefore $c' \in rR' = a'R' + d'R'$ because r was chosen to generate $a'R' + d'R'$, and this means that $c \in aR$ because $d'R' = \mathrm{Ker}(f)$. Hence b is an exact annihilator of a. □

The proof of Theorem 14.5 (ii) is quite technical, so we now illustrate how to construct a right exact annihilator of a matrix with entries in $\mathbf{Z}/m\mathbf{Z}$ for $m > 0$.

Example 14.6 Take $R = \mathbf{Z}/60\mathbf{Z}$ and consider the problem of constructing a right exact annihilator of

$$A = \begin{bmatrix} 7 & 2 \\ 9 & 6 \end{bmatrix}. \tag{14.8}$$

The first step is to reduce A to a diagonal matrix by finding invertible matrices $P, Q \in R^{2\times 2}$ with PAQ diagonal. The matrices

$$P = \begin{bmatrix} 1 & 0 \\ 33 & 1 \end{bmatrix} \quad \text{and} \quad Q = \begin{bmatrix} 43 & 34 \\ 0 & 1 \end{bmatrix} \tag{14.9}$$

have inverses

$$P^{-1} = \begin{bmatrix} 1 & 0 \\ 27 & 1 \end{bmatrix} \quad \text{and} \quad Q^{-1} = \begin{bmatrix} 7 & 2 \\ 0 & 1 \end{bmatrix}, \tag{14.10}$$

and it is easily verified that

$$PAQ = \begin{bmatrix} 1 & 0 \\ 0 & 12 \end{bmatrix}. \tag{14.11}$$

The next step is to construct a right exact annihilator of PAQ, and to do this we need to find exact annihilators of 1 and 12 in R. An exact annihilator of $a \in R$

can be found by dividing 60 by any generator of the ideal $a\mathbf{Z} + 60\mathbf{Z}$. More precisely, $60/\gcd\{a, 60\}$ is an exact annihilator of a, where $\gcd\{a, 60\}$ denotes the greatest common divisor of a and 60 in $\mathbf{Z}$. Therefore $0 \equiv 60/1 \bmod 60$ is an exact annihilator of 1 in R, and similarly $5 \equiv 60/12 \bmod 60$ is an exact annihilator of 12 in R. The proof of Lemma 14.4 then tells us that

$$C = \begin{bmatrix} 0 & 0 \\ 0 & 5 \end{bmatrix} \quad \text{and} \quad QC = \begin{bmatrix} 0 & 50 \\ 0 & 5 \end{bmatrix} \tag{14.12}$$

are right exact annihilators of PAQ and A respectively.

Notice that in the case $R = \mathbf{Z}/m\mathbf{Z}$ for some $m > 0$, each R-submodule of $R^{1\times n}$ must be the row space of some matrix because $R^{1\times n}$ only has finitely many elements. Furthermore, Theorem 14.3 tells us that if X is any R-submodule of $R^{1\times n}$ then the product of the cardinalities of X and $X^\perp$ is m^n. We can verify this in Example 14.6: direct computation reveals that $\mathrm{Row}(A) \subseteq (\mathbf{Z}/60\mathbf{Z})^{1\times 2}$ has 300 elements, and it is easy to check that $\mathrm{Row}(A)^\perp = \mathrm{Col}(QC)$ has 12 elements, giving a product of $3600 = 60^2$.

The other examples of exact annihilator rings that we are interested in are Boolean rings. A ring R is a *Boolean ring* if each element is multiplicatively idempotent, that is, if $aa = a$ for all $a \in R$. It follows from this definition that if R is a Boolean ring then R is commutative and satisfies $-a = a$ for each $a \in R$. It is also clear that $a = aaa$ for all $a \in R$, and consequently each finitely generated ideal of R is principal (see Lam [57, Theorem 4.23]). The powerset of any set can be given the structure of a Boolean ring by taking addition to be the symmetric difference of subsets and multiplication to be the intersection of subsets. Jacobson [45, section 8.5] provides a detailed discussion of this construction, which results in the field $\mathbf{Z}/2\mathbf{Z}$ in the simplest non-trivial case.[1]

Theorem 14.7 *If R is a Boolean ring then*

(i) *R is a commutative elementary divisor ring; and*

(ii) *R is an exact annihilator ring.*

Proof (i). Nicholson [69] calls a ring *clean* if each element is the sum of a unit and an idempotent, so since $a - 1$ is idempotent for each $a \in R$, it is obvious that R is

[1] It is easy to show that $\mathbf{Z}/2\mathbf{Z}$ is the only field which is also a Boolean ring.

clean. McGovern [66, Theorems 3.7 and 3.13; Proposition 3.14] has shown that if a clean ring has the property that each finitely generated ideal is principal then the ring must be an elementary divisor ring. Therefore R is a commutative elementary divisor ring because in addition to being clean, R is commutative and has all finitely generated ideals principal (see above).

(ii). As in the proof of Theorem 14.5 (ii), by (i) and Lemma 14.4 it is sufficient to show that each $a \in R$ has an exact annihilator. Since a is idempotent we have $a(1-a) = a - a = 0$, and as such $1-a$ is an annihilator of a. To show that $1-a$ is an exact annihilator of a, let $c \in R$ and suppose that $(1-a)c = 0$. We then want $c \in aR$, but this is obvious because $c = ac$. □

If R is a Boolean ring then the proof of Theorem 14.7 (ii) tells us how to construct a right exact annihilator of any given $A \in R^{m\times n}$ using the fact that $1-a = 1+a$ is an exact annihilator of $a \in R$. The following example gives some indication of the general form of a right exact annihilator of A.

Example 14.8 Let R be a Boolean ring and take

$$A = \begin{bmatrix} a & b \end{bmatrix} \tag{14.13}$$

for any $a, b \in R$. To construct a right exact annihilator of A we first need to find an invertible matrix $Q \in R^{2\times 2}$ with AQ diagonal. (Since A only has one row we can implicitly take $P = 1$.) The matrix

$$Q = \begin{bmatrix} 1 & b \\ 1+a & 1+b+ab \end{bmatrix} \tag{14.14}$$

has inverse

$$Q^{-1} = \begin{bmatrix} 1+b+ab & b \\ 1+a & 1 \end{bmatrix}, \tag{14.15}$$

and we have

$$AQ = \begin{bmatrix} a+b+ab & b+b+ab+ab \end{bmatrix} = \begin{bmatrix} a+b+ab & 0 \end{bmatrix} \tag{14.16}$$

because $b+b = 0$ and $ab+ab = 0$.

Next we need to construct a right exact annihilator of AQ, but before we can do this we must add a zero row to make AQ square diagonal. The resulting matrix

has diagonal entries $a + b + ab$ and 0, and these elements have exact annihilators $1 + a + b + ab$ and 1 respectively. The proof of Lemma 14.4 then tells us that

$$C = \begin{bmatrix} 1 + a + b + ab & 0 \\ 0 & 1 \end{bmatrix} \tag{14.17}$$

and

$$QC = \begin{bmatrix} 1 + a + b + ab & b \\ 1 + a + b + ab & 1 + b + ab \end{bmatrix} \tag{14.18}$$

are right exact annihilators of AQ and A respectively.

Theorem 14.7 (ii) tells us that every Boolean ring is an exact annihilator ring, and consequently every Boolean ring is exact by Proposition 13.7. This result is already known for some Boolean rings, as the powerset of any set is known to be self-injective when viewed as a ring (see Lam [56, Corollary 3.11D]). However, this alternative approach cannot be used to show that every Boolean ring is exact because not every Boolean ring arises as the powerset of a set. In fact, there even exist Boolean rings which are not self-injective, and so the observation that every Boolean ring is exact is of genuine interest. The standard construction of a non-self-injective Boolean ring is described by Lambek [58, page 45]: take the set of finite or cofinite subsets of any infinite set, with addition and multiplication given by symmetric difference and intersection respectively.

Residuated structures

15 Preliminary order theory

The semirings we will consider in the final sections of this thesis are all 'residuated' in the sense that division is possible up to a partial order. Specifically, an element a of such a 1-semiring S may or may not have a multiplicative inverse $a^{-1} \in S$ satisfying $aa^{-1} = 1$, but there will at least be some $b \in S$ satisfying $ab \leq 1$ in a suitable partial order. Moreover, there will always be a uniquely identifiable (in fact, a maximum) such element b, which we think of as being what remains after "dividing" 1 by a. As we will show in sections 18 and 19, the ability to divide in this way is very powerful, and it allows us to say a lot about linear algebra over S.

The above notion of residuation is actually applicable in a wider context—namely that of an ordered monoid acting on a poset—so before discussing residuation for semirings we will introduce this more general form of residuation (see section 17). In addition to providing a natural setting for residuation, actions of ordered monoids generalise actions of ordered semirings, i.e., ordered modules, and give us a sensible way to define what we mean by an 'anti-isomorphism' of ordered modules. We will study actions of ordered monoids in section 16, but first we recall some basic definitions and results of order theory.

The fundamental objects in order theory are partially ordered sets, or 'posets' for short. A *poset* is a set X together with a partial order $\leq$ on X, where a *partial order* is a binary relation on X which is *reflexive* ($x \leq x$ for all $x \in X$), *antisymmetric* ($x = y$ whenever $x, y \in X$ with $x \leq y$ and $y \leq x$) and *transitive* ($x \leq z$ whenever $x, y, z \in X$ with $x \leq y$ and $y \leq z$). A *join-semilattice* is a poset X in which each pair $x, y \in X$ has a *join*, or 'least upper bound', $x \vee y$ satisfying

$$x \vee y \leq z \quad \Leftrightarrow \quad x \leq z \text{ and } y \leq z \tag{15.1}$$

for all $z \in X$. Similarly, X is called a *meet-semilattice* if each pair $x, y \in X$ has a *meet*, or 'greatest lower bound', $x \wedge y$ satisfying

$$z \leq x \wedge y \quad \Leftrightarrow \quad z \leq x \text{ and } z \leq y \tag{15.2}$$

for all $z \in X$. A poset which is both a join-semilattice and a meet-semilattice is called a *lattice*. Lattices are an extremely important type of poset because they can be treated as algebraic objects (join and meet are commutative, associative binary operations on X). Birkhoff [8] covers the general theory of lattices in great detail and discusses some of its applications to other areas of mathematics.

The prototypical example of a lattice is the powerset of any set. In this case the partial order is given by subset inclusion and the join and meet operations are given by union and intersection respectively. Such a lattice X is special because the join and meet operations distribute over one another. Additionally, X has *bottom* and *top* elements $\bot \in X$ and $\top \in X$ respectively, satisfying $\bot \leq x \leq \top$ for all $x \in X$, and an involution $\overline{}$ called the *complement* which satisfies $x \vee \overline{x} = \top$ and $x \wedge \overline{x} = \bot$ for all $x \in X$. It follows automatically from the various axioms that the complement interchanges $\bot$ and $\top$. A lattice having these features is called a *Boolean algebra*, in honour of Boole [11]; the precise axioms were first formulated by Whitehead [85]. Note that there is a correspondence between Boolean algebras and the Boolean rings considered in section 14 (see Birkhoff [8, Theorem 10.3]).

Order-preserving functions between posets are called 'monotone', or sometimes 'isotone'. Specifically, a function $f\colon X \to Y$ between posets X and Y is *monotone* if it satisfies

$$x \leq y \quad \Rightarrow \quad fx \leq fy \tag{15.3}$$

for all $x, y \in X$. Monotone functions do not necessarily respect joins or meets if they exist in X and Y. That is, if X and Y are lattices then a monotone function $f\colon X \to Y$ may or may not satisfy $f(x \vee y) = fx \vee fy$ and $f(x \wedge y) = fx \wedge fy$ for all $x, y \in X$. If f satisfies the former condition then it is called *join-preserving*; if it satisfies the latter condition then it is called *meet-preserving*.

A bijection $f\colon X \to Y$ between posets X and Y is called an *order isomorphism* if f and f^{-1} are monotone. If X and Y happen to be lattices then f and f^{-1} are both join-preserving and meet-preserving (see Proposition 15.5, below). It is necessary to define an order isomorphism in this way because, as with continuous

functions in topology, the inverse of a monotone function need not be monotone. The requirement that f^{-1} is also monotone means that we have

$$x \leq y \quad \Leftrightarrow \quad fx \leq fy \tag{15.4}$$

for all $x, y \in X$. Any function $f\colon X \to Y$ satisfying (15.4) for all $x, y \in X$ is called an *order embedding*, and as such every order isomorphism is an order embedding. In fact, an order isomorphism is precisely a surjective order embedding because if $f\colon X \to Y$ is an order embedding then f is injective, by (15.4), so is an order isomorphism as a function $X \to \mathrm{Im}(f)$.

A function $f\colon X \to Y$ between posets X and Y is called *antitone* if it satisfies

$$x \leq y \quad \Rightarrow \quad fy \leq fx \tag{15.5}$$

for all $x, y \in X$. If f is a bijection and f^{-1} is also antitone then f is called an *order anti-isomorphism*. Order anti-isomorphisms do not preserve joins and meets if they exist, rather, an order anti-isomorphism $f\colon X \to Y$ satisfies $f(x \vee y) = fx \wedge fy$ and $f(x \wedge y) = fx \vee fy$ for all $x, y \in X$ if X and Y are lattices. Note that if X is a Boolean algebra then the complement $\overline{}\colon X \to X$ is an order anti-isomorphism.

Definition 15.1 Let X and Y be posets. A *Galois connection* between X and Y is a pair of functions $f\colon X \to Y$ and $g\colon Y \to X$ satisfying

$$y \leq fx \quad \Leftrightarrow \quad x \leq gy \tag{15.6}$$

for all $x \in X$ and all $y \in Y$.

If $f\colon X \to Y$ is an order anti-isomorphism between posets X and Y then f and f^{-1} constitute a Galois connection between X and Y. In general, Galois connections are weaker than order anti-isomorphisms, but this is not a bad thing because it means that Galois connections can arise in a wider variety of situations. Moreover, as we saw in Propositions 8.5 and 13.2, when a Galois connection does arise we can still deduce a great deal about the relationship between the two posets involved.

Proposition 15.2 *Let $f\colon X \to Y$ and $g\colon X \to Y$ be a Galois connection between posets X and Y. Then*

(i) *f and g are antitone;*

(ii) *$g \circ f$ and $f \circ g$ are monotone;*

(iii) *$x \leq (g \circ f)x$ for all $x \in X$ and $y \leq (f \circ g)y$ for all $y \in Y$; and*

(iv) *$f \circ g \circ f = f$ and $g \circ f \circ g = g$.*

Proof See Galatos et al. [27, Lemma 3.7]. □

In view of Proposition 15.2, a Galois connection between posets X and Y is sometimes defined to be a pair of antitone functions $f\colon X \to Y$ and $g\colon Y \to X$ satisfying $x \leq (g \circ f)x$ for all $x \in X$ and $y \leq (f \circ g)y$ for all $y \in Y$ (see Blyth [9, page 14]). This definition implies the condition in Definition 15.1 because

$$y \leq fx \quad \Rightarrow \quad x \leq (g \circ f)x \leq gy \tag{15.7}$$

and

$$x \leq gy \quad \Rightarrow \quad y \leq (f \circ g)y \leq fx \tag{15.8}$$

for all $x \in X$ and all $y \in Y$, that is, (15.6) holds for all $x \in X$ and all $y \in Y$.

Definition 15.3 Let X and Y be posets. An *adjunction* between X and Y is a pair of functions $f\colon X \to Y$ and $g\colon Y \to X$ satisfying

$$fx \leq y \quad \Leftrightarrow \quad x \leq gy \tag{15.9}$$

for all $x \in X$ and all $y \in Y$.

As we can see from Proposition 15.4, below, an adjunction is just the monotone version of a Galois connection. However, there is a slight subtlety in that the functions f and g in an adjunction do not satisfy identical conditions, whereas the functions in a Galois connection are essentially interchangeable. Specifically, if $f\colon X \to Y$ and $g\colon Y \to X$ constitute an adjunction then, by Proposition 15.4 (iii), $g \circ f$ satisfies $x \leq (g \circ f)x$ for all $x \in X$ but $f \circ g$ satisfies $(f \circ g)y \leq y$ for all $y \in Y$. Therefore $f \circ g$ and $g \circ f$ behave differently, and this means that an adjunction has an inherent direction.

To account for the fact that the two functions in Definition 15.3 have different properties, we will call f the *lower adjoint* and g the *upper adjoint*. This terminology is adapted from category theory, where f would be called the 'left adjoint' and g

would be called the 'right adjoint' (see Simmons [77, Example 1.3.3]).[1] It turns out that only one adjoint is needed to completely specify an adjunction; we may simply say that a function $f\colon X \to Y$ is a lower adjoint—or has an upper adjoint—and the function $g\colon Y \to X$ will be uniquely determinied by $gy = \max\{x \in X : fx \leq y\}$ for all $y \in Y$ (see Galatos et al. [27, Lemma 3.3]).

Proposition 15.4 *Let $f\colon X \to Y$ be a function between posets X and Y. If f has an upper adjoint $g\colon Y \to X$ then*

(i) *f and g are monotone;*

(ii) *$g \circ f$ and $f \circ g$ are monotone;*

(iii) *$x \leq (g \circ f)x$ for all $x \in X$ and $(f \circ g)y \leq y$ for all $y \in Y$; and*

(iv) *$f \circ g \circ f = f$ and $g \circ f \circ g = g$.*

Proof See Galatos et al. [27, Lemmas 3.1 to 3.3]. □

The lower adjoint of an adjunction between posets is often called a 'residuated' function, after Blyth and Janowitz [10]. However, we prefer to reserve this term for the case of certain multiplication functions being lower adjoints (see Definition 17.1) because it is suggestive of "division". Note that Blyth and Janowitz [10, page 11] define a monotone function $f\colon X \to Y$ to be residuated if there is a monotone function $g\colon Y \to X$ satisfying $x \leq (g \circ f)x$ for all $x \in X$ and $(f \circ g)y \leq y$ for all $y \in Y$. By a similar argument to the one given above for Galois connections, this definition is equivalent to the definition of a lower adjoint.

If $f\colon X \to Y$ is an order isomorphism between posets X and Y then f is a lower adjoint, with upper adjoint f^{-1}. It is also an upper adjoint, with lower adjoint f^{-1}, so in this case there is an adjunction between X and Y in both directions. In general, though, the existence of an adjunction in one direction between X and Y does not imply that there will be an adjunction the other direction (put another way, not every lower adjoint is also an upper adjoint). Adjunctions are weaker than order isomorphisms, but are still strong enough to automatically preserve additional structure that X and Y might have. The following result exemplifies this feature of adjunctions, and allows us to immediately see that an order isomorphism must be both join-preserving and meet-preserving.

[1] It is safer to use the terms 'lower' and 'upper' here because 'left' and 'right' adjoints could be too easily confused with left and right actions in section 17.

Proposition 15.5 *Let $f\colon X \to Y$ be a function between lattices X and Y. If f has an upper adjoint $g\colon Y \to X$ then f is join-preserving and g is meet-preserving.*

Proof See Galatos et al. [27, Lemma 3.5]. □

A *closure operator* on a poset X is a monotone function $h\colon X \to X$ which is *expanding* ($x \leq hx$ for all $x \in X$) and idempotent. It is clear from Proposition 15.4 that if $f\colon X \to Y$ is a lower adjoint between posets X and Y, with upper adjoint $g\colon Y \to X$, then $g \circ f$ is a closure operator on X. Notice that $g \circ f$ is idempotent because we have $(g \circ f) \circ (g \circ f) = g \circ (f \circ g \circ f) = g \circ f$ by Proposition 15.4 (iv). If f and g were a Galois connection between X and Y then Proposition 15.2 would tell us that $f \circ g$ is also a closure operator, but the fact that f and g have different roles in an adjunction means that $f \circ g$ is not a closure operator on Y. Instead, $f \circ g$ is an *interior operator* on Y because it is monotone, idempotent and satisfies $(f \circ g)y \leq y$ for all $y \in Y$.

16 Ordered monoids and actions

The semirings we will study in sections 18 to 20 are partially ordered, so underlying each one is an ordered (multiplicative) monoid. In this section we recall the definition of an ordered monoid and we consider actions of ordered monoids on posets. By characterising isomorphisms in this setting, we arrive at a natural definition of an anti-isomorphism of ordered algebraic structures. In section 17 we will introduce residuation in the context of ordered monoids acting on posets.

An *ordered monoid* is a monoid $(M, \cdot, 1)$ together with a partial order $\leq$ on M that is compatible with multiplication in the sense that

$$a \leq b \quad \Rightarrow \quad ca \leq cb \text{ and } ac \leq bc \tag{16.1}$$

for all $a, b, c \in M$. In other words, if a monoid M is to be an ordered monoid then the multiplication functions $c-\colon M \to M$ and $-c\colon M \to M$ must be monotone for each $c \in M$. Any monoid M can be made into an ordered monoid by setting $a \leq b$ if and only if $a = b$, but, of course, we are primarily interested in less trivial examples. The motivating examples to have in mind are the group $(\mathbf{R}, +, 0)$ with the standard ordering of $\mathbf{R}$ and the Boolean monoid $(\{0, 1\}, \min, 1)$ with $0 < 1$.[1] Notice that

[1] The Boolean monoid is also known as the 'two-element semilattice'.

these monoids are the multiplicative monoids of the finitary tropical semiring **FT** and the Boolean semiring **B** respectively.

Recall from section 2 that if M is a monoid then a right monoid action of M on set X is a function $\cdot\colon X \times M \to X$ satisfying $x(ab) = (xa)b$ and $x1 = x$ for all $a, b \in M$ and all $x \in X$. This definition is completely algebraic, so is not sufficient for considering actions of M when M is an ordered monoid and X is a poset. In such a case, the action of M on X ought to be compatible with the partial orders on M and X in much the same way that multiplication and the partial order on M are compatible. That is, all the actions of M we consider should satisfy some version of (16.1). The following definition makes this requirement precise.

Definition 16.1 Let M be an ordered monoid, let X be a poset and let $\cdot$ be a right monoid action of M on X. Then $(X, \cdot)$ is a *right M-poset* if

$$a \leq b \quad \Rightarrow \quad xa \leq xb \tag{16.2}$$

and

$$x \leq y \quad \Rightarrow \quad xa \leq ya \tag{16.3}$$

for all $a, b \in M$ and all $x, y \in X$.

Condition (16.2) says that if a poset X is to be a right M-poset then for each $x \in X$ the function $x-\colon M \to X$ must be monotone. Similarly, (16.3) requires the function $-a\colon X \to X$ to be monotone for each $a \in M$. These two conditions are sometimes combined into an equivalent single requirement: the right action of M on X must be monotone as a function $X \times M \to X$ (see Bulman-Fleming and Mahmoudi [15, page 443]), where the partial order on $X \times M$ is given by $(x, a) \leq (y, b)$ if and only if $x \leq y$ and $a \leq b$. However, for reasons that will become clear in section 17, we prefer to treat the action of M on X as two separate families of monotone functions, parametrised by the elements of X and M respectively.

A right M-poset is an ordered algebraic structure, that is, it is partially ordered and has some algebraic structure (the right action of M) which interacts with the partial order. Structure preserving functions between M-posets should therefore respect the partial orders and the actions of M. Specifically, a structure preserving function $f\colon X \to Y$ between right M-posets X and Y should be monotone and should satisfy $f(xa) = (fx)a$ for all $a \in M$ and all $x \in X$. If such a function is an order isomorphism then it is called an *isomorphism* of right M-posets. Note that,

as usual in algebra, the inverse of an isomorphism f satisfies $f^{-1}(ya) = \big(f^{-1}y\big)a$ for all $a \in M$ and all $y \in Y$, so is automatically a structure preserving function in the appropriate sense.

We would also like to consider anti-isomorphisms between M-posets, but to do this we obviously need a suitable notion of what an 'anti-isomorphism' even is in this context. The above definition of an isomorphism is not much help in this regard, as it turns out to be incorrect to naively define an anti-isomorphism of M-posets to be an order anti-isomorphism that respects the actions of M.[1] Instead, we need to take inspiration from a more sophisticated characterisation of isomorphisms. The following definition is the first step in this direction.

Definition 16.2 Let M be an ordered monoid and let $f\colon X \to Y$ be a function between right M-posets X and Y. Then f is *right M-monotone* if

$$xa \leq y \quad \Rightarrow \quad (fx)a \leq fy \tag{16.4}$$

for all $a \in M$ and all $x, y \in X$.

By taking $a = 1$ in Definition 16.2, we see that a right M-monotone function $f\colon X \to Y$ is monotone (this works because the action of M on a right M-poset is assumed to be a monoid action), and by taking $y = xa$ we see that f satisfies $(fx)a \leq f(xa)$ for all $a \in M$ and all $x \in X$. Therefore a right M-monotone function is almost a structure preserving function in the sense described above. Definition 16.2 does not guarantee that $f(xa) \leq (fx)a$ for all $a \in M$ and all $x \in X$ though, so in general the notion of a right M-monotone function is weaker than that of a structure preserving function. However, as the following result shows, if f happens to have an inverse that is also right M-monotone then it is true that $f(xa) \leq (fx)a$ for all $a \in M$ and all $x \in X$ as well.

Proposition 16.3 *Let M be an ordered monoid and let $f\colon X \to Y$ be a bijection between right M-sets X and Y. If f and f^{-1} are right M-monotone then we have $f(xa) = (fx)a$ for all $a \in M$ and all $x \in X$.*

Proof Let $a \in M$ and let $x \in X$. Then since f is right M-monotone we have $(fx)a \leq f(xa)$ by taking $y = xa$ in Definition 16.2. It therefore remains to show

[1] The complement anti-isomorphism on a Boolean algebra does not satisfy $\overline{b \wedge a} = \overline{b} \wedge a$ for all elements a, b.

that $f(xa) \leq (fx)a$. Since f^{-1} is also right M-monotone we have

$$xa = \big(f^{-1}(fx)\big)a \leq f^{-1}((fx)a), \tag{16.5}$$

because $(fx)a \leq (fx)a$, and thus $f(xa) \leq (fx)a$ because f is monotone. □

Proposition 16.3 tells us that if $f\colon X \to Y$ is right M-monotone and has a right M-monotone inverse then f is an order isomorphism satisfying $f(xa) = (fx)a$ for all $a \in M$ and all $x \in X$, i.e., f is an isomorphism of right M-posets. Conversely, a monotone function $f\colon X \to Y$ satisfying $f(xa) = (fx)a$ for all $a \in M$ and all $x \in X$ is clearly right M-monotone, so since the inverse of a right M-poset isomorphism $f\colon X \to Y$ is also monotone with $f^{-1}(ya) = \big(f^{-1}y\big)a$ for all $a \in M$ and all $y \in Y$ it follows that f and f^{-1} are both right M-monotone. We may therefore characterise an isomorphism of right M-posets as a right M-monotone function that has a right M-monotone inverse.

Although it does not initially appear to be useful, being able to characterise isomorphisms of M-posets in terms of M-monotone functions is certainly progress, as the definition of an M-monotone function accounts for the partial orders and the actions of M in one combined condition. This might seem counter-productive (it is impossible to say how an M-monotone function interacts with the actions of M without mentioning partial orders), but in fact it is just what we need because it captures the correct sense in which an anti-isomorphism of M-posets should interact with the actions of M. Specifically, an anti-isomorphism should satisfy an antitone version of (16.4). The two obvious candidates for such a condition are

$$xa \leq y \quad \Rightarrow \quad fy \leq (fx)a \tag{16.6}$$

and

$$xa \leq y \quad \Rightarrow \quad (fy)a \leq fx. \tag{16.7}$$

Condition (16.7) is much better than (16.6) because it would enable us to show that the composition of two anti-isomorphisms is an isomorphism, but for technical reasons (16.7) is still not quite right. In section 19 we will see that when an anti-isomorphism arises naturally it is between a right M-poset and a left M-poset, not between two right M-posets. We should therefore make one of the actions in (16.7) a left action instead of a right action.

Definition 16.4 Let M be an ordered monoid and let $f\colon X \to Y$ be a function from a right M-poset X to a left M-poset Y. Then f is *M-antitone* if

$$xa \leq y \quad \Rightarrow \quad a(fy) \leq fx \tag{16.8}$$

for all $a \in M$ and all $x, y \in X$.

Note that a *left M-poset* is just a poset Y together with a left monoid action of M, such that for each $y \in Y$ and each $a \in M$ the functions $-y\colon M \to Y$ and $a-\colon Y \to Y$ are monotone.

We now define an *anti-isomorphism* of M-posets to be an M-antitone bijection $f\colon X \to Y$ whose inverse is also M-antitone. If we assume that X denotes a right M-poset and that Y denotes a left M-poset then, technically, Definition 16.4 cannot be applied to f^{-1} because f^{-1} is a function from a left (not a right) M-poset to a right (not a left) M-poset. The upshot of this is that we must actually use a slightly modified version of Definition 16.4 for f^{-1}; when we say that $f^{-1}\colon Y \to X$ is *M-antitone* we really mean that f^{-1} satisfies

$$ax \leq y \quad \Rightarrow \quad \big(f^{-1}y\big)a \leq f^{-1}x \tag{16.9}$$

for all $a \in M$ and all $x, y \in Y$.

To prevent ambiguity, we will exclusively refer to the function from the right M-poset to the left M-poset as being the anti-isomorphism. This convention means that an anti-isomorphism always satisfies (16.8), and that the inverse of an anti-isomorphism always satisfies (16.9). If there is an anti-isomorphism from a right M-poset X to a left M-poset Y then we will say that X and Y are anti-isomorphic (without regard for the direction of the anti-isomorphism) and we will write $X \mathrel{\widetilde{\approx}} Y$.

As the name suggests, an M-antitone function from a right M-poset to a left M-poset is antitone (take $a = 1$ in Definition 16.4). An anti-isomorphism from a right M-poset to a left M-poset is therefore an order anti-isomorphism, and thus the composition of (the inverse of) an anti-isomorphism with an anti-isomorphism is an order isomorphism. In fact, as a consequence of the similarity between M-monotone and M-antitone functions, such a composition is actually an isomorphism of M-posets. This result confirms that our notion of anti-isomorphism of M-posets is sensible.

Proposition 16.5 *Let M be an ordered monoid. Then the composition of the inverse of an anti-isomorphism of M-posets with an anti-isomorphism of M-posets is an isomorphism of right M-posets.*

Proof Let $f\colon X \to Z$ be an anti-isomorphism from a right M-poset X to a left M-poset Z and let $g\colon Y \to Z$ be an anti-isomorphism from a right M-poset Y to Z. We want to show that $g^{-1} \circ f$ is an isomorphism of right M-posets, so in view of Proposition 16.3 it is sufficient to show that $g^{-1} \circ f$ and $\big(g^{-1} \circ f\big)^{-1}$ are right M-monotone. Since f and g are anti-isomorphisms we have

$$xa \leq y \quad \Rightarrow \quad a(fy) \leq fx \quad \Rightarrow \quad \big(g^{-1}(fx)\big)a \leq g^{-1}(fy) \tag{16.10}$$

for all $a \in M$ and all $x, y \in X$, by (16.8) and (16.9), and as such $g^{-1} \circ f$ is right M-monotone by Definition 16.2. A similar argument shows that $\big(g^{-1} \circ f\big)^{-1} = f^{-1} \circ g$ is right M-monotone. Hence $g^{-1} \circ f$ is an isomorphism of right M-posets. □

17 Residuation and enriched categories

In this section we introduce residuation for M-posets, where throughout M denotes an ordered monoid. A residuated M-poset turns out to have the structure of an enriched category (see Proposition 17.3), so we also investigate the conditions under which an enriched category gives rise to a residuated M-poset. We begin by defining what it means for an M-poset to be residuated.

Definition 17.1 Let M be an ordered monoid. A right M-poset X is *residuated* if there is a function $d_{\mathrm{R}}\colon X \times X \to M$, called *right residuation*, satisfying

$$xa \leq y \quad \Leftrightarrow \quad a \leq d_{\mathrm{R}}(x, y) \tag{17.1}$$

for all $a \in M$ and all $x, y \in X$.

Equivalently, by Definition 15.3, a right M-poset X is residuated if and only if for each $x \in X$ the function $x-\colon M \to X$ has an upper adjoint $d_{\mathrm{R}}(x, -)\colon X \to M$. Notice that each upper adjoint $d_{\mathrm{R}}(x, -)$ satisfies $a \leq d_{\mathrm{R}}(x, xa)$ and $x d_{\mathrm{R}}(x, y) \leq y$ for all $a \in M$ and all $y \in X$. This suggests that $d_{\mathrm{R}}(x, -)$ can be thought of as approximating "division by x", and accordingly $d_{\mathrm{R}}(x, y)$ is often written as $x \setminus y$ (see Galatos et al. [27, page 92]) or $y \mathbin{.\!\cdot} x$ (see Blyth and Janowitz [10, page 211]).

However, in view of Proposition 17.3, below, we prefer to think of $d_{\mathrm{R}}(x, y)$ as being a kind of distance from x to y.

A left M-poset X is *residuated* if there is a function $d_{\mathrm{L}} \colon X \times X \to M$, called *left residuation*, satisfying

$$ax \leq y \quad \Leftrightarrow \quad a \leq d_{\mathrm{L}}(y, x) \tag{17.2}$$

for all $a \in M$ and all $x, y \in X$. As above, this definition says that for each $x \in X$ the function $-x \colon M \to X$ is a lower adjoint, with upper adjoint $d_{\mathrm{L}}(-, x) \colon X \to M$ satisfying $a \leq d_{\mathrm{L}}(ax, x)$ and $d_{\mathrm{L}}(y, x)x \leq y$ for all $a \in M$ and all $y \in X$. Note that it is fairly standard to write $y \,/\, x$ instead of $d_{\mathrm{L}}(y, x)$.

Definition 17.2 A *residuated monoid* is an ordered monoid M that is residuated as a right M-poset and as a left M-poset.

An ordered monoid M is a right M-poset and a left M-poset via multiplication, so Definition 17.2 tells us that if M is a residuated monoid then there are functions $d_{\mathrm{R}}, d_{\mathrm{L}} \colon M \times M \to M$ satisfying

$$a \leq d_{\mathrm{L}}(c, b) \quad \Leftrightarrow \quad ab \leq c \quad \Leftrightarrow \quad b \leq d_{\mathrm{R}}(a, c) \tag{17.3}$$

for all $a, b, c \in M$. This condition says that the functions $d_{\mathrm{R}}(a, -) \colon M \to M$ and $d_{\mathrm{L}}(-, b) \colon M \to M$ are upper adjoints for the multiplication functions $a- \colon M \to M$ and $-b \colon M \to M$ respectively, so by Proposition 15.4 (i) each such multiplication function is monotone. The definition of an ordered monoid requires these functions to be monotone anyway, so we do not appear to have gained anything by applying Proposition 15.4 (i). However, this observation is useful because it means that when verifying Definition 17.2 we do not actually need to check that we have an ordered monoid.

If the poset underlying a residuated monoid M is a lattice then M is called a *residuated lattice*. In such a case, Proposition 15.5 tells us that for each $c \in M$ the multiplication functions $c- \colon M \to M$ and $-c \colon M \to M$ are join-preserving because they are lower adjoints. That is, we have

$$c(a \vee b) = ca \vee cb \tag{17.4}$$

and

$$(a \vee b)c = ac \vee bc \tag{17.5}$$

for all $a, b, c \in M$. Note that if M was not residuated, but was just an ordered monoid that happened to be a lattice, then (17.4) and (17.5) would not necessarily hold because monotone functions need not be join-preserving. Proposition 15.5 also tells us that for each $c \in M$ the functions $d_{\mathrm{R}}(c, -)\colon M \to M$ and $d_{\mathrm{L}}(-, c)\colon M \to M$ are meet-preserving because they are upper adjoints. Therefore

$$d_{\mathrm{R}}(c, a \wedge b) = d_{\mathrm{R}}(c, a) \wedge d_{\mathrm{R}}(c, b) \tag{17.6}$$

and

$$d_{\mathrm{L}}(a \wedge b, c) = d_{\mathrm{L}}(a, c) \wedge d_{\mathrm{L}}(b, c) \tag{17.7}$$

for all $a, b, c \in M$.

A Boolean algebra can be viewed as an ordered monoid by taking meet as multiplication and the top element as 1, and it turns out that every such ordered monoid $(M, \wedge, \top)$ is a residuated lattice with $d_{\mathrm{R}}(a, b) = \overline{a} \vee b$ for all $a, b \in M$ (see Galatos et al. [27, Lemma 3.22]). In particular, the Boolean monoid $(\{0, 1\}, \min, 1)$ is a residuated lattice with $d_{\mathrm{R}}(a, b) = \max\{1 - a, b\}$ for all $a, b \in \{0, 1\}$. Since the join operation on a lattice is always commutative, every Boolean algebra is a commutative monoid, and this means that we have $d_{\mathrm{L}}(b, a) = d_{\mathrm{R}}(a, b)$ for all elements a, b of a Boolean algebra. Not every residuated monoid is commutative though, so in general right and left residuation will be different.

Another important class of residuated monoids is the *ordered groups*, that is, the ordered monoids which happen to be groups. If G is an ordered group then we have

$$a \leq cb^{-1} \quad \Leftrightarrow \quad ab \leq c \quad \Leftrightarrow \quad b \leq a^{-1}c \tag{17.8}$$

for all $a, b, c \in G$, and as such G is a residuated monoid with $d_{\mathrm{R}}(a, c) = a^{-1}c$ and $d_{\mathrm{L}}(c, b) = cb^{-1}$ for all $a, b, c \in G$. In particular, the group $(\mathbf{R}, +, 0)$ is a residuated monoid with $d_{\mathrm{R}}(a, b) = d_{\mathrm{L}}(b, a) = b - a$ for all $a, b \in \mathbf{R}$. This ordered group is actually a lattice, with $a \vee b = \max\{a, b\}$ and $a \wedge b = \min\{a, b\}$ for all $a, b \in \mathbf{R}$, and thus $(\mathbf{R}, +, 0)$ is a residuated lattice. We will give a few more examples of residuated lattices in sections 19 and 20; Galatos et al. [27, section 3.4] give several more.

Proposition 17.3 *Let M be an ordered monoid and let X be a residuated right M-poset. Then*

(i) $d_{\mathrm{R}}(x, y) \cdot d_{\mathrm{R}}(y, z) \leq d_{\mathrm{R}}(x, z)$ *for all $x, y, z \in X$; and*

(ii) $1 \leq d_{\mathrm{R}}(x,x)$ *for all* $x \in X$.

Proof (i). Let $x, y, z \in X$ and take $a = d_{\mathrm{R}}(x,y)$. Then in particular $a \leq d_{\mathrm{R}}(x,y)$, so $xa \leq y$ by (17.1). Similarly, if we take $b = d_{\mathrm{R}}(y,z)$ then we have $yb \leq z$, and thus $x(ab) = (xa)b \leq yb \leq z$ because X is a right M-poset. Therefore $ab \leq d_{\mathrm{R}}(x,z)$ by (17.1) again, which means that $d_{\mathrm{R}}(x,y) \cdot d_{\mathrm{R}}(y,z) \leq d_{\mathrm{R}}(x,z)$.

(ii). Let $x \in X$. Then since the right action of M on X is a monoid action we have $x1 = x$. Hence $1 \leq d_{\mathrm{R}}(x,x)$ by (17.1) because $x1 \leq x$. □

Proposition 17.3 tells us that if M is an ordered monoid then every residuated right M-poset can be interpreted as a small category enriched over M. The general definition of an enriched category is rather complicated (see Kelly [50, section 1.2]), so we give a simplified definition which only applies in this setting. Let $(M, \cdot, 1)$ be an ordered monoid. Then an *M-category*, or a category *enriched* over M, is a set X together with a function $d\colon X \times X \to M$ satisfying $d(x,y) \cdot d(y,z) \leq d(x,z)$ and $1 \leq d(x,x)$ for all $x, y, z \in X$, i.e., d must satisfy Proposition 17.3 (i) and (ii). The objects in an M-category are the elements of the set X, just as in a normal small category, but instead of having (sets of) morphisms between objects we associate a monoid element $d(a,b)$ with each pair $a, b \in X$ of objects. An enriched category can also be thought of as a heavily generalised metric space: if we take M to be the monoid $(\mathbf{R}, +, 0)$ with the reverse ordering of $\mathbf{R}$ then an M-category is a set X together with a function $d\colon X \times X \to \mathbf{R}$ satisfying $d(x,z) \leq d(x,y) + d(y,z)$ and $d(x,x) \leq 0$ for all $x, y, z \in X$ (see Lawvere [59]).

If M is an ordered monoid then every residuated right M-poset is an M-category, but does the converse hold? This seems unlikely in general, as there are at least two obstructions to an arbitrary M-category (X,d) being a residuated right M-poset. Firstly, we can define a natural *preorder* (a reflexive and transitive binary relation) on X using d, but this preorder is not necessarily a partial order. If we are willing to quotient X by an equivalence relation though, this preorder turns into a partial order (see Simmons [77, Exercise 2.6.7]). Secondly, there is no obvious way to define a right action of M on X, and so we essentially have to just assume that one exists. The following result therefore appears to be the best partial converse we can obtain in general.

Theorem 17.4 *Let M be an ordered monoid and let (X,d) be an M-category. If*

there is a function $\cdot\colon X \times M \to X$ *satisfying*

$$1 \leq d(xa, y) \quad \Leftrightarrow \quad a \leq d(x, y) \tag{17.9}$$

and

$$b \leq d(xa, x(ab)) \tag{17.10}$$

for all $a, b \in M$ *and all* $x, y \in X$ *then there is an equivalence relation* $\sim$ *on* X *with* $X/\!\sim$ *a residuated right* M*-poset and* $d_{\mathrm{R}}([x], [y]) = d(x, y)$ *for all* $x, y \in X$.

Proof The binary relation $\preceq$ on X defined by

$$x \preceq y \quad \Leftrightarrow \quad 1 \leq d(x, y) \tag{17.11}$$

is reflexive because $1 \leq d(x, x)$ for all $x \in X$. This relation is also transitive because if $x, y, z \in X$ with $1 \leq d(x, y)$ and $1 \leq d(y, z)$ then the assumption that M is an ordered monoid gives $1 \leq d(x, y) \cdot d(y, z) \leq d(x, z)$. Therefore $\preceq$ is a preorder on X, and as such the binary relation $\sim$ on X defined by

$$x \sim y \quad \Leftrightarrow \quad x \preceq y \text{ and } y \preceq x \tag{17.12}$$

is an equivalence relation on X. Moreover, the quotient $X/\!\sim$ is a poset with

$$[x] \leq [y] \quad \Leftrightarrow \quad x \preceq y \tag{17.13}$$

for all $x, y \in X$ (see Simmons [77, Exercise 2.6.7]).

To show that $X/\!\sim$ is a residuated right M-poset, we first need to define a right monoid action of M on $X/\!\sim$. We would like to define $[x]a = [xa]$ for all $a \in M$ and all $x \in X$, but it is not yet clear that this is well-defined. That is, we need to check that $x \sim y$ implies $xa \sim ya$ for all $a \in M$ and all $x, y \in X$. In view of (17.12) it is sufficient to show that $x \preceq y$ implies that $xa \preceq ya$ for all $a \in M$ and all $x, y \in X$.

Let $a \in M$, let $x, y \in X$ and suppose that $x \preceq y$. Then $1 \leq d(x, y)$ by (17.11), so $d(y, ya) \leq d(x, y) \cdot d(y, ya) \leq d(x, ya)$ because M is an ordered monoid. Now since $1 \leq d(ya, ya)$ we have $a \leq d(y, ya)$ by (17.9), and thus $a \leq d(y, ya) \leq d(x, ya)$. Therefore $1 \leq d(xa, ya)$ by (17.9) again, which means that $xa \preceq ya$. As described above, this result allows us to define $[x]a = [xa]$ for all $a \in M$ and all $x \in X$. By

combining this definition with (17.9), (17.11) and (17.13) we then obtain

$$[x]a \leq [y] \quad \Leftrightarrow \quad a \leq d(x, y) \tag{17.14}$$

for all $a \in M$ and all $x, y \in X$.

If we can show that $X/{\sim}$ is a right M-poset then (17.14) will ensure that it is residuated, so it remains to show that our proposed right action of M on $X/{\sim}$ is actually a right monoid action that is monotone in the sense of Definition 16.1. However, (17.14) already tells us that for each $x \in X$ the function $[x]{-}\colon M \to X/{\sim}$ is monotone (because it is a lower adjoint), and when we showed above that $x \preceq y$ implies $xa \preceq ya$ for all $a \in M$ and all $x, y \in X$ we essentially showed that each function $-a\colon X/{\sim} \to X/{\sim}$ is also monotone. Therefore $X/{\sim}$ is a residuated right M-poset, provided we really do have a right monoid action.

To show that we have a right monoid action of M on $X/{\sim}$, we first need to show that $[x](ab) = ([x]a)b$ for all $a, b \in M$ and all $x \in X$. One inequality is clear, as we have

$$([x]a)b = [xa]b \leq [x(ab)] = [x](ab) \tag{17.15}$$

by (17.10) and (17.14). For the reverse inequality, (17.14) gives $a \leq d(x, xa)$ and $b \leq d(xa, (xa)b)$ because $[x]a \leq [xa]$ and $[xa]b \leq [(xa)b]$, so

$$ab \leq d(x, xa) \cdot d(xa, (xa)b) \leq d(x, (xa)b) \tag{17.16}$$

because M is an ordered monoid. Therefore $[x](ab) \leq [(xa)b] = ([x]a)b$ by (17.14) again, and as such $[x](ab) = ([x]a)b$.

Finally, we need to show that $[x]1 = [x]$ for all $x \in X$. Since $1 \leq d(x, x)$, (17.14) gives $[x]1 \leq [x]$. For the reverse inequality we have $1 \leq d(x, x1)$ by (17.9) because $1 \leq d(x1, x1)$, so $x \preceq x1$ by (17.11). This means that $[x] \leq [x1] = [x]1$, and thus $[x]1 = [x]$. Therefore $X/{\sim}$ is a right M-poset. Hence (17.14) confirms that $X/{\sim}$ is a residuated right M-poset with $d_{\mathrm{R}}([x], [y]) = d(x, y)$ for all $x, y \in X$. □

Linear algebra over residuated lattices

18 Matrix residuation

Recall from section 17 that a residuated lattice is an ordered algebraic structure that is simultaneously a lattice and a (multiplicative) monoid. By taking addition to be the lattice join operation we can treat a residuated lattice as a semiring, and this makes it possible to work with matrices over residuated lattices. In this section we show that the behaviour of matrices with entries in a residuated lattice is controlled by left and right residuation, but it turns out that residuation alone is not powerful enough to answer all of our questions about matrices. Specifically, we can use residuation to describe kernel classes (see Proposition 18.5), but a residuated lattice may or may not be exact.

In section 19 we will consider special residuated lattices where residuation can be expressed in terms of an involution (in fact, a conjugation), and we will show that the presence of such an involution guarantees exactness. It is not known whether every exact residuated lattice must be 'involutive' in this sense however, but in section 20 we will show that at least some non-involutive residuated lattices are also not exact. Our main new result (Corollary 20.5) is that the residuated lattice of subsets of a finite monoid is exact if and only if the monoid is a group.

Let $(M, \cdot, 1)$ be a residuated lattice. Since M is a lattice, join is a commutative, associative binary operation on M, and as such $(M, \vee)$ is a commutative semigroup. Furthermore, $(M, \cdot)$ is a semigroup satisfying $c(a \vee b) = ca \vee cb$ and $(a \vee b)c = ac \vee bc$ for all $a, b, c \in M$, by (17.4) and (17.5). This means that Definition 4.1 (i) holds for $(M, \vee, \cdot)$, and so to make M a semiring we just need to produce local identities that satisfy Definition 4.1 (ii). That is, for each non-empty finite $L \subseteq M$ we need to find $0_L, 1_L \in M$ with $a \vee b0_L = a1_L = a$ and $a \vee 0_L b = 1_L a = a$ for all $a, b \in L$. A residuated lattice is assumed to be a monoid, so it is obvious that we can simply take $1_L = 1$, but the existence of 0_L is less clear.

Proposition 18.1 *Let M be a residuated lattice and let $L \subseteq M$. If L is non-empty and finite then there is some $0_L \in M$ satisfying $a \vee b0_L = a$ and $a \vee 0_L b = a$ for all $a, b \in L$.*

Proof Since L is non-empty and finite we may take

$$0_L = \bigwedge_{a,b \in L} d_{\mathrm{R}}(b, a) \wedge d_{\mathrm{L}}(a, b), \tag{18.1}$$

where d_{R} and d_{L} denote right and left residuation respectively. Now let $a, b \in L$. Then $0_L \leq d_{\mathrm{R}}(b, a)$ by (15.2), and so (17.3) gives $b0_L \leq a$. Therefore $a \vee b0_L = a$ by (15.1). A dual argument involving $d_{\mathrm{L}}(a, b)$ confirms that $a \vee 0_L b = a$. □

Proposition 18.1 completes the above verification of Definition 4.1 (ii) in the case of a residuated lattice, and consequently any residuated lattice $(M, \cdot, 1)$ can be viewed as a semiring $S = (M, \vee, \cdot)$. In fact, S is a 1-semiring because M has an identity element. The finitary tropical semiring $\mathbf{FT} = (\mathbf{R}, \max, +)$ and the Boolean semiring $\mathbf{B} = (\{0, 1\}, \max, \min)$ arise in this way, as $(\mathbf{R}, +, 0)$ and $(\{0, 1\}, \min, 1)$ are residuated lattices with $\vee = \max$ (see page 103). Notice that if $L \subseteq \mathbf{FT}$ is non-empty and finite then (18.1) tells us that we can take $0_L = \min\{a - b : a, b \in L\}$ because $d_{\mathrm{R}}(b, a) = d_{\mathrm{L}}(a, b) = a - b$ for all $a, b \in \mathbf{FT}$.

Now let S be a residuated lattice. Since S is a semiring we can add and multiply matrices over S in accordance with (6.1) and (6.2), and (because of the way we are treating S as a semiring) these operations interact with the partial order on S in several important ways. Firstly, the sum of two matrices is precisely their entrywise join because addition on S is the join operation. This means that if we extend the partial order on S to $S^{m \times n}$ by setting $A \leq B$ if and only if $A_{ij} \leq B_{ij}$ for all $1 \leq i \leq m$ and all $1 \leq j \leq n$, then each $S^{m \times n}$ is a lattice and we have $A + B = A \vee B$ for all $A, B \in S^{m \times n}$. Note that the meet of A and B in $S^{m \times n}$ is similarly given by the entrywise meet of A and B.

Matrix multiplication also interacts with the extended partial orders on each $S^{m \times n}$: if $A, B \in S^{m \times n}$ with $A \leq B$ then we have $CA \leq CB$ for all $C \in S^{p \times m}$ and $AD \leq BD$ for all $D \in S^{n \times q}$. Put another way, for each $C \in S^{p \times m}$ and each $D \in S^{n \times q}$ the functions $C-\colon S^{m \times n} \to S^{p \times n}$ and $-D\colon S^{m \times n} \to S^{m \times q}$ are monotone. In particular, the right (monoid) action of S on $S^{m \times 1}$ satisfies

$$a \leq b \quad \Rightarrow \quad xa \leq xb \tag{18.2}$$

and

$$x \leq y \quad \Rightarrow \quad xa \leq ya \tag{18.3}$$

for all $a, b \in S$ and all $x, y \in S^{m\times 1}$, and as such each $S^{m\times 1}$ is a right S-poset by Definition 16.1. In fact, the following result shows that each $S^{m\times 1}$ is a residuated right S-poset.

Proposition 18.2 *If S is a residuated lattice then d_{R} and d_{L} extend to functions $d_{\mathrm{R}}\colon S^{m\times n} \times S^{m\times q} \to S^{n\times q}$ and $d_{\mathrm{L}}\colon S^{m\times q} \times S^{n\times q} \to S^{m\times n}$ satisfying*

$$A \leq d_{\mathrm{L}}(C, B) \quad \Leftrightarrow \quad AB \leq C \quad \Leftrightarrow \quad B \leq d_{\mathrm{R}}(A, C) \tag{18.4}$$

for all $A \in S^{m\times n}$, all $B \in S^{n\times q}$ and all $C \in S^{m\times q}$.

Proof Let $A \in S^{m\times n}$, let $B \in S^{n\times q}$ and let $C \in S^{m\times q}$. Since $AB \leq C$ if and only if $(AB)_{ik} \leq C_{ik}$ for all $1 \leq i \leq m$ and all $1 \leq k \leq q$, (6.2) gives

$$AB \leq C \quad \Leftrightarrow \quad \bigvee_{j=1}^{n} A_{ij}B_{jk} \leq C_{ik} \text{ for all } i,\, k. \tag{18.5}$$

Therefore

$$AB \leq C \quad \Leftrightarrow \quad A_{ij}B_{jk} \leq C_{ik} \text{ for all } i,\, j,\, k \tag{18.6}$$

by (15.1), and thus

$$AB \leq C \quad \Leftrightarrow \quad B_{jk} \leq d_{\mathrm{R}}(A_{ij}, C_{ik}) \text{ for all } i,\, j,\, k \tag{18.7}$$

by (17.3). Finally, (15.2) gives

$$AB \leq C \quad \Leftrightarrow \quad B_{jk} \leq \bigwedge_{i=1}^{m} d_{\mathrm{R}}(A_{ij}, C_{ik}) \text{ for all } j,\, k, \tag{18.8}$$

and as such $AB \leq C$ if and only if $B \leq d_{\mathrm{R}}(A, C)$ once we define $d_{\mathrm{R}}(A, C) \in S^{n\times q}$ by

$$(d_{\mathrm{R}}(A, C))_{jk} = \bigwedge_{i=1}^{m} d_{\mathrm{R}}(A_{ij}, C_{ik}) \tag{18.9}$$

for all $1 \leq j \leq n$ and all $1 \leq k \leq q$. Notice that this definition recovers right residuation on S in the case $m = n = q = 1$.

A dual argument confirms that $AB \leq C$ if and only if $A \leq d_{\mathrm{L}}(C, B)$, where $d_{\mathrm{L}}(C, B) \in S^{m \times n}$ is defined by

$$(d_{\mathrm{L}}(C, B))_{ij} = \bigwedge_{k=1}^{q} d_{\mathrm{L}}(C_{ik}, B_{jk}) \tag{18.10}$$

for all $1 \leq i \leq m$ and all $1 \leq j \leq n$. □

In the case $n = q = 1$, Proposition 18.2 tells us that if S is a residuated lattice then there is a function $d_{\mathrm{R}} \colon S^{m \times 1} \times S^{m \times 1} \to S$ satisfying

$$xa \leq y \quad \Leftrightarrow \quad a \leq d_{\mathrm{R}}(x, y) \tag{18.11}$$

for all $a \in S$ and all $x, y \in S^{m \times 1}$. The existence of such a function makes $S^{m \times 1}$ a residuated right S-poset (see Definition 17.1). Note that Hollings and Kambites [40] write $\langle x \mid y \rangle$ instead of $d_{\mathrm{R}}(x, y)$ in the case S is the completed tropical semiring $\overline{\mathbf{T}} = (\mathbf{R} \cup \{-\infty, \infty\}, \max, +)$.

If S is a residuated lattice and X is a right S-module then we have

$$x + x = x1 + x1 = x(1 + 1) = x(1 \vee 1) = x1 = x \tag{18.12}$$

for all $x \in X$, and as such addition on X is idempotent. This means that we can partially order X by setting $x \leq y$ if and only if $x + y = y$. Notice that in the case $X = S^{m \times 1}$ this partial order is precisely the entrywise partial order defined above because addition on $S^{m \times 1}$ is entrywise join. Definition 4.4 (i) then ensures that every right S-module is a right S-poset. Moreover, if $f \colon X \to Y$ is a right S-linear function between right S-modules X and Y then f preserves the S-poset structures of X and Y, i.e., f is monotone and satisfies $f(xa) = (fx)a$ for all $a \in S$ and all $x \in X$. Every isomorphism of right S-modules is therefore an isomorphism of right S-posets.

Our use of addition to partially order a given S-module means that joins always exist in S-modules: if X is an S-module and $x, y \in X$ then $x + y$ satisfies

$$\begin{aligned} x + y \leq z \quad &\Leftrightarrow \quad x + y + z = z \\ &\Leftrightarrow \quad x + z = z \text{ and } y + z = z \\ &\Leftrightarrow \quad x \leq z \text{ and } y \leq z \end{aligned} \tag{18.13}$$

for all $z \in X$ because addition on X is idempotent. So, since the join operation on an S-module is given by addition, every join-preserving function—hence every order isomorphism—between S-modules satisfies Definition 5.1 (i). Every isomorphism of right S-posets is therefore an isomorphism of right S-modules, and thus the notions of right S-module isomorphism and right S-poset isomorphism are equivalent. In view of Proposition 16.5, this suggests that S-poset anti-isomorphism is the strongest notion of anti-isomorphism between S-modules that we could hope for.

As we discussed briefly above, some S-modules not only have joins, but are in fact lattices. For instance, each $S^{m\times 1}$ is a lattice with entrywise join and meet. If $A \in S^{m\times n}$ then the column space of A is a right S-submodule of $S^{m\times 1}$, so is closed under joins in $S^{m\times 1}$, but it is not clear whether $\mathrm{Col}(A)$ is closed under meets in $S^{m\times 1}$. Example 18.4, below, demonstrates that in general the column space of a matrix is not closed under meets. However, the following result shows that we can always use residuation to define (possibly different) meets in $\mathrm{Col}(A)$.

Proposition 18.3 *Let S be a residuated lattice and let $A \in S^{m\times n}$. Then $\mathrm{Col}(A)$ is a lattice, with the meet of $x, y \in \mathrm{Col}(A)$ given by $Ad_{\mathrm{R}}(A, x \wedge y)$.*

Proof This follows from a more general result, which we first prove. Let $f\colon X \to Y$ be a lower adjoint between lattices X and Y, with upper adjoint $g\colon Y \to X$. We will show that $\mathrm{Im}(f)$ has meets, with the meet of fx and fy given by $(f \circ g)(fx \wedge fy)$ for all $x, y \in X$. Let $x, y, z \in X$. Then

$$fz \leq (f \circ g)(fx \wedge fy) \quad \Leftrightarrow \quad z \leq (g \circ f \circ g)(fx \wedge fy) \tag{18.14}$$

by Definition 15.3, where $(g \circ f \circ g)(fx \wedge fy) = g(fx \wedge fy)$ by Proposition 15.4 (iv). Therefore

$$fz \leq (f \circ g)(fx \wedge fy) \quad \Leftrightarrow \quad fz \leq fx \wedge fy \tag{18.15}$$

by Definition 15.3 again, and thus

$$fz \leq (f \circ g)(fx \wedge fy) \quad \Leftrightarrow \quad fz \leq fx \text{ and } fz \leq fy \tag{18.16}$$

by (15.2). This means that $(f \circ g)(fx \wedge fy)$ is the meet of fx and fy in $\mathrm{Im}(f)$, as claimed.

Now, Proposition 18.2 tells us that the function $A-\colon S^{n\times 1} \to S^{m\times 1}$ is a lower adjoint, with upper adjoint $d_{\mathrm{R}}(A,-)\colon S^{m\times 1} \to S^{n\times 1}$. Hence $\mathrm{Col}(A) = \mathrm{Im}(A-)$ has meets, with the meet of Av and Aw in $\mathrm{Col}(A)$ given by $Ad_{\mathrm{R}}(A, Av \wedge Aw)$ for all $v, w \in S^{n\times 1}$. □

Example 18.4 The column space of

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{bmatrix} \in \mathbf{B}^{3\times 2} \tag{18.17}$$

comprises $\left[\begin{smallmatrix} 0 \\ 0 \\ 0 \end{smallmatrix}\right]$, $\left[\begin{smallmatrix} 1 \\ 0 \\ 1 \end{smallmatrix}\right]$, $\left[\begin{smallmatrix} 0 \\ 1 \\ 1 \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} 1 \\ 1 \\ 1 \end{smallmatrix}\right]$, so is not closed under meets in $\mathbf{B}^{3\times 1}$ because

$$\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \wedge \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \notin \mathrm{Col}(A). \tag{18.18}$$

Instead, Proposition 18.3 tells us that the meet of $\left[\begin{smallmatrix} 1 \\ 0 \\ 1 \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} 0 \\ 1 \\ 1 \end{smallmatrix}\right]$ in $\mathrm{Col}(A)$ is $\left[\begin{smallmatrix} 0 \\ 0 \\ 0 \end{smallmatrix}\right]$ because $d_{\mathrm{R}}\Big(A, \left[\begin{smallmatrix} 0 \\ 0 \\ 1 \end{smallmatrix}\right]\Big) = \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right]$ and $A\left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right] = \left[\begin{smallmatrix} 0 \\ 0 \\ 0 \end{smallmatrix}\right]$.

Notice that if $A \in S^{m\times n}$ is a matrix with entries in a residuated lattice S then the function $Ad_{\mathrm{R}}(A,-)\colon S^{m\times 1} \to S^{m\times 1}$ used in Proposition 18.3 is an interior operator on $S^{m\times 1}$ because it is the lower adjoint of an adjunction composed with the upper adjoint (see page 96). Alternatively, we could use Proposition 18.2 to show that $Ad_{\mathrm{R}}(A,-)$ is monotone, idempotent and satisfies $Ad_{\mathrm{R}}(A,y) \leq y$ for all $y \in S^{m\times 1}$. Cohen et al. [20, page 403] call $Ad_{\mathrm{R}}(A,-)$ the 'canonical projector' onto $\mathrm{Col}(A)$, and, as they note, $Ad_{\mathrm{R}}(A,-)$ fixes $\mathrm{Col}(A)$ pointwise. This property can be seen by appealing to Proposition 15.4 (iv), or by using the direct argument given in the proof of Proposition 18.5 (ii) below.

Similarly, the function $d_{\mathrm{R}}(A, A-)\colon S^{n\times 1} \to S^{n\times 1}$ is a closure operator on $S^{n\times 1}$ because it is monotone, idempotent and satisfies $v \leq d_{\mathrm{R}}(A, Av)$ for all $v \in S^{n\times 1}$. The first part of the following result shows that the kernel of $\mathrm{Row}(A)$ can be characterised as the pairs of vectors which have the same closure under $d_{\mathrm{R}}(A, A-)$; $\operatorname{Ker}\mathrm{Row}(A)$ is the set-theoretic kernel of this closure operator.

Proposition 18.5 *Let S be a residuated lattice and let $A \in S^{m\times n}$. Then*

(i) $\operatorname{Ker}\operatorname{Row}(A) = \big\{(v, v') \in S^{n\times 1} \times S^{n\times 1} : d_{\mathrm{R}}(A, Av) = d_{\mathrm{R}}(A, Av')\big\}$; *and*

(ii) $\operatorname{Ker}\operatorname{Row}(A) = \operatorname{Ker}^2(F)$, *where* $F = \big\{(v, d_{\mathrm{R}}(A, Av)) : v \in S^{n\times 1}\big\}$.

Proof (i). Let $v, v' \in S^{n\times 1}$ and suppose that $(v, v') \in \operatorname{Ker}\operatorname{Row}(A)$. Then $Av = Av'$ by Proposition 8.6 (i), and thus $d_{\mathrm{R}}(A, Av) = d_{\mathrm{R}}(A, Av')$. Conversely, suppose that $d_{\mathrm{R}}(A, Av) = d_{\mathrm{R}}(A, Av')$. Since $Av \leq Av$, Proposition 18.2 gives $v \leq d_{\mathrm{R}}(A, Av)$. Therefore $v \leq d_{\mathrm{R}}(A, Av')$ because $d_{\mathrm{R}}(A, Av) = d_{\mathrm{R}}(A, Av')$, and thus $Av \leq Av'$ by Proposition 18.2 again. A similar argument shows that $Av' \leq Av$, and as such $Av = Av'$. Hence $(v, v') \in \operatorname{Ker}\operatorname{Row}(A)$ by Proposition 8.6 (i).

(ii). Let $v, v' \in S^{n\times 1}$ and suppose that $(v, v') \in \operatorname{Ker}\operatorname{Row}(A)$. Then (i) gives $d_{\mathrm{R}}(A, Av) = d_{\mathrm{R}}(A, Av')$. By Definition 8.1, to show that $(v, v') \in \operatorname{Ker}^2(F)$ we need to show that $yv = yv'$ for all $y \in \operatorname{Ker}(F)$, so let $y \in \operatorname{Ker}(F)$. Then $yv = yd_{\mathrm{R}}(A, Av)$ and $yv' = yd_{\mathrm{R}}(A, Av')$ by Definition 8.3 and the definition of F. Therefore $yv = yv'$, as required for $(v, v') \in \operatorname{Ker}^2(F)$, because $d_{\mathrm{R}}(A, Av) = d_{\mathrm{R}}(A, Av')$.

Conversely, suppose that $(v, v') \in \operatorname{Ker}^2(F)$. Then $yv = yv'$ for all $y \in \operatorname{Ker}(F)$ by Definition 8.1. Now let $w \in S^{n\times 1}$. Then since $Aw \leq Aw$ we have $w \leq d_{\mathrm{R}}(A, Aw)$ by Proposition 18.2, and thus $Aw \leq Ad_{\mathrm{R}}(A, Aw)$ because multiplication by a fixed matrix is a monotone function. Since $d_{\mathrm{R}}(A, Aw) \leq d_{\mathrm{R}}(A, Aw)$, Proposition 18.2 also gives $Ad_{\mathrm{R}}(A, Aw) \leq Aw$, and as such $Aw = Ad_{\mathrm{R}}(A, Aw)$. We therefore have $uAw = uAd_{\mathrm{R}}(A, Aw)$ for all $u \in S^{1\times m}$ and all $w \in S^{n\times 1}$, so, by Definition 8.3 and the definition of F, each $uA \in \operatorname{Ker}(F)$. This means that $uAv = uAv'$ for all $u \in S^{1\times m}$. Hence $(v, v') \in \operatorname{Ker}\operatorname{Row}(A)$ by Definition 8.1. □

Part (ii) of Proposition 18.5 tells us that the closure operator $d_{\mathrm{R}}(A, A-)$ captures the essential information about $\operatorname{Ker}\operatorname{Row}(A)$, in the sense that no information is lost if $\operatorname{Ker}\operatorname{Row}(A)$ is reduced to just the pairs $(v, d_{\mathrm{R}}(A, Av))$ for $v \in S^{n\times 1}$. To compute the classes of $\operatorname{Ker}\operatorname{Row}(A)$ using d_{R}, start with the closed vectors, i.e., the vectors $d_{\mathrm{R}}(A, x)$ for $x \in \operatorname{Col}(A)$. Each such vector will be the maximum element in its class, and the class of $d_{\mathrm{R}}(A, x)$ will comprise as many of the vectors beneath $d_{\mathrm{R}}(A, x)$ as possible. Specifically, the class of $d_{\mathrm{R}}(A, x)$ will be the set of all $v \in S^{n\times 1}$ which satisfy

$$x \leq y \quad \Leftrightarrow \quad v \leq d_{\mathrm{R}}(A, y) \tag{18.19}$$

for all $y \in \operatorname{Col}(A)$, as these are precisely the vectors whose closure is $d_{\mathrm{R}}(A, x)$.

19 Involutive residuated lattices

In section 18 we explained how matrix residuation can be used to describe kernels over residuated lattices, thus addressing the first of our main problems. Our second main problem is to decide whether an arbitrary semiring is exact, but in the case of residuated lattices we do not yet have any tools with which to do this. Since a residuated lattice may or may not be exact (see Corollary 20.5), we must assume that our residuated lattices have some additional structure if we are to decide whether they are exact. In this section we argue that a very useful piece of extra structure to insist upon is a certain kind of involution, and we show that every residuated lattice equipped with such an involution is exact. We also show that the involution induces an anti-isomorphism between the row space and column space of each matrix.

Definition 19.1 A residuated lattice S is *involutive* if there is an involution $\overline{}$ on S satisfying $\overline{a} = d_{\mathrm{R}}\big(a, \overline{1}\big) = d_{\mathrm{L}}\big(\overline{1}, a\big)$ for all $a \in S$ (see Wille [87, page 38]).

Combining Definition 19.1 with (17.3), we see that the involution on an involutive residuated lattice S must satisfy

$$a \leq \overline{b} \quad \Leftrightarrow \quad ab \leq \overline{1} \quad \Leftrightarrow \quad b \leq \overline{a} \tag{19.1}$$

for all $a, b \in S$, and thus in cases where $\overline{1} = 1$ the involution can be thought of as approximating a "multiplicative inverse". For instance, every ordered group $(G, \cdot, 1)$ whose underlying poset is a lattice is an involutive residuated lattice with $\overline{a} = a^{-1}$ for all $a \in G$. However, this analogy breaks down for more exotic involutive residuated lattices: if $(S, \vee, \wedge)$ is a Boolean algebra then S is an involutive residuated lattice because the complement $\overline{}\colon S \to S$ satisfies (19.1), with $1 = \top$, but the complement is the worst imaginable approximation to a multiplicative inverse because it satisfies $a \wedge \overline{a} = \bot$ for all $a \in S$.

In accordance with our convention (see page 63), we extend the involution on an involutive residuated lattice S to matrices by setting $\overline{A}_{ji} = \overline{A_{ij}}$ for all $A \in S^{m \times n}$.

Proposition 19.2 *If S is an involutive residuated lattice then*

$$AB \leq C \quad \Rightarrow \quad B\overline{C} \leq \overline{A} \tag{19.2}$$

for all $A \in S^{m \times n}$, all $B \in S^{n \times q}$ and all $C \in S^{m \times q}$.

Proof Let $A \in S^{m\times n}$, let $B \in S^{n\times q}$ and let $C \in S^{m\times q}$. Then (as in the proof of Proposition 18.2) we have

$$AB \leq C \quad \Leftrightarrow \quad A_{ij}B_{jk} \leq C_{ik} \text{ for all } i,\, j,\, k \tag{19.3}$$

and

$$B\overline{C} \leq \overline{A} \quad \Leftrightarrow \quad B_{jk}\overline{C}_{ki} \leq \overline{A}_{ji} \text{ for all } i,\, j,\, k \tag{19.4}$$

$$\Leftrightarrow \quad B_{jk}\overline{C_{ik}} \leq \overline{A_{ij}} \text{ for all } i,\, j,\, k. \tag{19.5}$$

It therefore suffices to show that

$$ab \leq c \quad \Rightarrow \quad b\overline{c} \leq \overline{a} \tag{19.6}$$

for all $a, b, c \in S$, so let $a, b, c \in S$ and suppose that $ab \leq c$. Then $ab \leq \overline{\overline{c}}$ because $\overline{}$ is an involution, and thus $ab\overline{c} \leq \overline{1}$ by (19.1). Hence $b\overline{c} \leq \overline{a}$ by (19.1) again, as required. □

In particular, Proposition 19.2 tells us that if S is an involutive residuated lattice then the function $\overline{}\colon S^{m\times 1} \to S^{1\times m}$ is S-antitone in the sense of Definition 16.4. Moreover, the inverse function $\overline{}\colon S^{1\times m} \to S^{m\times 1}$ is also S-antitone (in the dual sense) because a double application of Proposition 19.2 gives

$$ax \leq y \quad \Rightarrow \quad x\overline{y} \leq \overline{a} \quad \Rightarrow \quad \overline{y}a \leq \overline{x} \tag{19.7}$$

for all $a \in S$ and all $x, y \in S^{1\times m}$. The involution on S therefore induces an anti-isomorphism of S-posets from $S^{m\times 1}$ to $S^{1\times m}$. That is, $S^{m\times 1} \eqsim S^{1\times m}$ as S-posets. Note that these anti-isomorphisms exist even if S is not commutative, whereas S needs to be commutative for the transpose operation to induce an isomorphism of S-posets between $S^{m\times 1}$ and $S^{1\times m}$.

Proposition 19.2 will feature heavily in all of our manipulations of matrices over involutive residuated lattices, so it is worth thinking about how to remember what it allows us to do. Intuitively, (19.2) says that the three terms in an inequality of the form $AB \leq C$ may be cycled one place to the left, at the cost of introducing a '$\overline{}$' to the two terms that cross the '$\leq$'. As we saw above, the fact that $\overline{}$ is an involution means that we can apply this procedure once more to obtain $\overline{C}A \leq \overline{B}$, and consequently the terms in an inequality of the form $AB \leq C$ may also be cycled one place to the right at the same cost. In the results that follow we make frequent

use of this cycling technique without further mention.

Note that a straightforward adaptation of the proof of Proposition 19.2 shows that if $A, B \in S^{m\times n}$ with $A \leq B$ then $\overline{B} \leq \overline{A}$.[1] Therefore $\overline{\ }$ is antitone as a function from $S^{m\times n}$ to $S^{n\times m}$. This result is also easy to remember, thanks to the above rule which says that any term crossing a '$\leq$' must pick up a '$\overline{\ }$', so we will use it too without further mention.[2]

If S is an involutive residuated lattice then we can express matrix residuation in terms of the involution on S. Specifically, Proposition 18.2 gives $d_{\mathrm{R}}(A, C) = \overline{\overline{C}A}$ for all $A \in S^{m\times n}$ and all $C \in S^{m\times q}$ because

$$AB \leq C \quad \Leftrightarrow \quad \overline{C}A \leq \overline{B} \quad \Leftrightarrow \quad B \leq \overline{\overline{C}A} \tag{19.8}$$

for all $B \in S^{n\times q}$. Dually, $d_{\mathrm{L}}(C, B) = \overline{B\overline{C}}$ for all $B \in S^{n\times q}$ and all $C \in S^{m\times q}$. These results illustrate why the presence of an involution lets us tackle problems (e.g., obtaining exactness) that residuation alone is not powerful enough for, even if the final answer does not actually refer to the involution: the involution gives us an "intermediate place" to work. When proving facts involving residuation we first apply the involution to turn column vectors into row vectors (for example), then do some manipulation of row vectors, and finally apply the involution a second time so that the end result can be phrased in terms of residuation. If a residuated lattice is not involutive then this middle step is not possible, and as a result we are limited in what we can prove using residuation.

Lemma 19.3, below, is a good example of a result that is ostensibly only about residuation, but that is actually proved using an involution in the way described above. Notice that Lemma 19.3 improves upon Proposition 18.5 (ii) for involutive residuated lattices because if $\mathrm{Row}(A) = \mathrm{Ker}(F)$ then $\mathrm{Ker}\,\mathrm{Row}(A) = \mathrm{Ker}^2(F)$.

Lemma 19.3 *Let S be an involutive residuated lattice and let $A \in S^{m\times n}$. Then* $\mathrm{Row}(A) = \mathrm{Ker}(F)$, *where* $F = \big\{(v, d_{\mathrm{R}}(A, Av)) : v \in S^{n\times 1}\big\}$.

Proof First recall from Definition 8.3 that

$$\mathrm{Ker}(F) = \big\{y \in S^{1\times n} : yv = yd_{\mathrm{R}}(A, Av) \text{ for all } v \in S^{n\times 1}\big\}. \tag{19.9}$$

[1] Or just apply Proposition 19.2 to $AI \leq B$, where I is an identity matrix local to A and $\overline{B}$.

[2] We can use these cycling rules to show that if S is an involutive residuated lattice then $ab \leq \overline{c}$ if and only if $a \leq \overline{bc}$ for all $a, b, c \in S$. This means that the category associated with S (create an object for each $a \in S$ and a morphism from $a \in S$ to $b \in S$ if $a \leq b$; see Awodey [3, pages 9–10]) is '$*$-autonomous' (see Barr [6]). Such categories are used to model linear logic (see Melliès [67]).

Since $Av = Ad_{\mathrm{R}}(A, Av)$ for all $v \in S^{n\times 1}$ (see page 112) we have $uAv = uAd_{\mathrm{R}}(A, Av)$ for all $u \in S^{1\times m}$ and all $v \in S^{n\times 1}$, and as such $x \in \mathrm{Ker}(F)$ for all $x \in \mathrm{Row}(A)$. It therefore remains to show that $\mathrm{Ker}(F) \subseteq \mathrm{Row}(A)$.

Now let $y \in \mathrm{Ker}(F)$. Then in particular $y\overline{y} = yd_{\mathrm{R}}(A, A\overline{y})$, where $y\overline{y} \leq \overline{1}$ because $1y \leq y$. Therefore

$$y\overline{A\overline{y}A} = yd_{\mathrm{R}}(A, A\overline{y}) \leq \overline{1} \tag{19.10}$$

because $d_{\mathrm{R}}(A, x) = \overline{\overline{x}A}$ for all $x \in S^{m\times 1}$ (see above), and thus $y \leq \overline{A\overline{y}}A$. We also have $\overline{A\overline{y}}A \leq y$ because $A\overline{y} \leq A\overline{y}$, so in fact $y = \overline{A\overline{y}}A$, and as such $y \in \mathrm{Row}(A)$. Hence $\mathrm{Ker}(F) \subseteq \mathrm{Row}(A)$. □

Theorem 19.4 *If S is an involutive residuated lattice then S is exact.*

Proof If $A \in S^{m\times n}$ then by Lemma 19.3 there is some $F \subseteq S^{n\times 1} \times S^{n\times 1}$ satisfying $\mathrm{Row}(A) = \mathrm{Ker}(F)$. Hence S is right exact by Lemma 9.12; a dual argument confirms that S is also left exact. □

The finitary tropical semiring $\mathbf{FT} = (\mathbf{R}, \max, +)$ is an involutive residuated lattice with $\overline{a} = -a$ for all $a \in \mathbf{FT}$ (see page 114), so is exact by Theorem 19.4. The completed tropical semiring $\overline{\mathbf{T}} = (\mathbf{R} \cup \{-\infty, \infty\}, \max, +)$ is also exact because it too is an involutive residuated lattice—although proving this requires some case analysis because of how $\overline{\mathbf{T}}$ is defined (see Definition 2.3). Note that the involution on $\overline{\mathbf{T}}$ extends the involution on $\mathbf{FT}$ by interchanging $-\infty$ and ∞. Every Boolean algebra is also an involutive residuated lattice, so in particular the Boolean semiring $\mathbf{B} = (\{0, 1\}, \max, \min)$ is exact. We will explore some consequences of $\mathbf{B}$ being exact in section 20.

In section 18 we demonstrated that if S is a residuated lattice then kernel classes can be described using residuation. Now, thanks to Theorem 19.4, we know that if in addition S is involutive then S is exact. Our third and final main problem is to understand the relationship between the row space and column space of each $A \in S^{m\times n}$, and, as we discussed in section 10, one way to do this is via a conjugation on S. The definition of a conjugation is practically rigged to ensure that if S is involutive then the involution on S is a conjugation, but since the involution on an involutive residuated lattice is not a standard involution, section 10 does not provide any insight into the properties of the induced bijection between $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$. It turns out that $\mathrm{Row}(A)$ and $\mathrm{Col}(A)$ are anti-isomorphic as S-posets.

Theorem 19.5 *If S is an involutive residuated lattice then*

(i) *the involution on S is a conjugation; and*

(ii) $\mathrm{Row}(A) \mathrel{\overline{\approx}} \mathrm{Col}(A)$ *as S-posets for all $A \in S^{m\times n}$.*

Proof (i). Let $A \in S^{m\times n}$. Then $\overline{A\overline{uA}}A = d_{\mathrm{L}}(uA, A)A = uA$ for all $u \in S^{1\times m}$ and $A\overline{\overline{Av}A} = Ad_{\mathrm{R}}(A, Av) = Av$ for all $v \in S^{n\times 1}$ (see pages 112 and 116). Therefore Definition 10.1 (i) and (ii) are satisfied with $M = N = A$, and as such the involution on S is a conjugation.

(ii). Continuing from above, we know that the function $\mathrm{Col}(A) \to \mathrm{Row}(A)$ given by $x \mapsto \overline{x}A$ is a bijection with inverse given by $x \mapsto A\overline{x}$. It is therefore sufficient to show that this function and its inverse are S-antitone in the appropriate senses. Firstly, we have

$$xa \leq y \quad \Rightarrow \quad a\overline{y} \leq \overline{x} \quad \Rightarrow \quad a\overline{y}A \leq \overline{x}A \tag{19.11}$$

for all $a \in S$ and all $x, y \in S^{m\times 1}$ because multiplication by a fixed matrix is a monotone function, so in particular the function $\mathrm{Col}(A) \to \mathrm{Row}(A)$ given by $x \mapsto \overline{x}A$ is S-antitone by Definition 16.4. Dually, the inverse of this function is also S-antitone because we have

$$ax \leq y \quad \Rightarrow \quad \overline{y}a \leq \overline{x} \quad \Rightarrow \quad A\overline{y}a \leq A\overline{x} \tag{19.12}$$

for all $x, y \in S^{1\times n}$. Hence the function $\mathrm{Col}(A) \to \mathrm{Row}(A)$ is an anti-isomorphism of S-posets. □

Hollings and Kambites [40, Theorem 2.4] have already shown that if S is $\mathbf{FT}$ or $\overline{\mathbf{T}}$ then $\mathrm{Row}(A) \mathrel{\overline{\approx}} \mathrm{Col}(A)$ for all $A \in S^{m\times n}$. Their notion of anti-isomorphism would not be suitable for all involutive residuated lattices, but it turns out to be equivalent to our notion of anti-isomorphism of S-posets in the case S is $\mathbf{FT}$ or $\overline{\mathbf{T}}$.

20 Subsets of groups and monoids

If S is an involutive residuated lattice then Theorem 19.4 tells us that S is an exact semiring. It then follows from Corollary 12.5 (and its dual) that the group semiring SG is exact for every finite group G, so in particular every finite group semiring $\mathbf{B}G$ is exact because $\mathbf{B}$ is an involutive residuated lattice. As we discussed in section 7,

in the case of $\mathbf{B}$ it is not actually necessary to restrict to finite groups, but since Corollary 12.5 applies to arbitrary exact semirings it can only be used to establish that $\mathbf{B}G$ is exact for finite groups G. In this section we show that $\mathbf{B}G$ is an involutive residuated lattice for every group G, and as a result every group semiring $\mathbf{B}G$ is in fact exact.

Let $(M, \cdot, 1)$ be a monoid, finite or infinite. Then the powerset of M can be made into an ordered monoid by setting $VW = \{st : s \in V \text{ and } t \in W\}$ for all $V, W \subseteq M$ and by taking the partial order to be subset inclusion. Moreover, this ordered monoid $(\mathrm{Pow}(M), \cdot, \{1\})$ turns out to be residuated, with

$$d_{\mathrm{R}}(V, W) = \{r \in M : Vr \subseteq W\} \tag{20.1}$$

and

$$d_{\mathrm{L}}(W, V) = \{r \in M : rV \subseteq W\} \tag{20.2}$$

for all $V, W \subseteq M$ (see Galatos et al. [27, section 3.4.10]). Therefore $(\mathrm{Pow}(M), \cdot, \{1\})$ is a residuated lattice, because $\mathrm{Pow}(M)$ is a lattice with join and meet given by union and intersection respectively.

Now recall from section 7 that if M is a monoid then the semiring $\mathbf{B}M$ can be identified with the powerset of M, and that the sum and product of $V, W \in \mathbf{B}M$ are given by $V \cup W$ and $\{st : s \in V \text{ and } t \in W\}$ respectively. When viewed in this way it is clear that $\mathbf{B}M$ is just the semiring corresponding to the residuated lattice $(\mathrm{Pow}(M), \cdot, \{1\})$, and consequently all the results in section 18 apply to $\mathbf{B}M$. However, the results in section 19 cannot be applied to $\mathbf{B}M$ unless M is a group.

Theorem 20.1 *Let M be a monoid. Then the residuated lattice $\mathbf{B}M$ is involutive if and only if M is a group.*

Proof ($\Rightarrow$). Suppose that $\mathbf{B}M$ is involutive. Then $\overline{\emptyset} = M$ and $\overline{M} = \emptyset$ because the involution $\overline{\ }$ on $\mathbf{B}M$ is an order anti-isomorphism. Now let $s \in M$. Then we have $s\overline{Ms} \subseteq \overline{M} = \emptyset$ because $Ms \subseteq Ms$, and thus $s\overline{Ms} = \emptyset$. If $\overline{Ms}$ were non-empty we would have $st \in s\overline{Ms}$ for some $t \in \overline{Ms}$, so since $s\overline{Ms} = \emptyset$ we must also have $\overline{Ms} = \emptyset$. Therefore $Ms = \overline{\emptyset} = M$, and as such there is some $r \in M$ with $rs = 1$. That is, s has a left inverse in M. Hence M is a group.

($\Leftarrow$). Suppose that M is a group. To show that $\mathbf{B}M$ is involutive we need to produce an involution $\overline{\ }$ on $\mathbf{B}M$ that satisfies $\overline{V} = d_{\mathrm{R}}\big(V, \overline{\{1\}}\big) = d_{\mathrm{L}}\big(\overline{\{1\}}, V\big)$ for all $V \in \mathbf{B}M$. The involution on $\mathbf{B}M$ defined by $\overline{V} = \big\{s^{-1} : s \in V\big\}$ for all $V \in \mathbf{B}M$

satisfies this condition because we have

$$d_{\mathrm{R}}\big(V, \overline{\{1\}}\big) = d_{\mathrm{R}}(V, \{1\}) = \{r \in M : sr = 1 \text{ for all } s \in V\} = \overline{V} \tag{20.3}$$

and

$$d_{\mathrm{R}}\big(\overline{\{1\}}, V\big) = d_{\mathrm{R}}(\{1\}, V) = \{r \in M : rs = 1 \text{ for all } s \in V\} = \overline{V} \tag{20.4}$$

for all $V \in \mathbf{B}M$, by (20.1) and (20.2) respectively. □

Corollary 20.2 *If G is a group then $\mathbf{B}G$ is exact.*

Proof By Theorem 20.1, $\mathbf{B}G$ is an involutive residuated lattice. Hence $\mathbf{B}G$ is exact by Theorem 19.4. □

Now we have established that $\mathbf{B}G$ is exact for every group G, we turn our attention to the converse problem: if M is a monoid, does exactness of $\mathbf{B}M$ force M to be a group? Theorem 20.1 tells us that if $\mathbf{B}M$ is an involutive residuated lattice then M must be a group, but this does not quite answer the question because $\mathbf{B}M$ could conceivably be exact without being involutive. The problem is easy to resolve for cancellative monoids, as the following result implies that if M is cancellative with $\mathbf{B}M$ exact then M must actually be a group.

Proposition 20.3 *Let M be a monoid and let $r \in M$. If $\mathbf{B}M$ is right exact and r is left cancellative then r has a left inverse in M.*

Proof Let $V, V' \in \mathbf{B}M$ and suppose that $(V, V') \in \operatorname{Ker} \operatorname{Row}(\{r\})$. Then $rV = rV'$, and as such $rv \in rV'$ for all $v \in V$. Therefore for each $v \in V$ there is some $v' \in V'$ with $rv = rv'$, but since r is left cancellative we have $v = v' \in V'$. This means that $V \subseteq V'$, so $V = V'$ by symmetry. Definition 8.3 then gives $\{1\} \in \operatorname{Ker}^2 \operatorname{Row}(\{r\})$, because $1V = 1V'$, and thus $\{1\} \in \operatorname{Row}(\{r\})$ by Proposition 9.10. Hence r has a left inverse in M because there is some $W \in \mathbf{B}M$ satisfying $\{1\} = Wr$. □

Proposition 20.3 deals with cancellative monoids, but the above problem remains open for arbitrary monoids; must M be a group if $\mathbf{B}M$ is exact, or does there exist a non-cancellative monoid M with $\mathbf{B}M$ exact? However, if we restrict to finite monoids then the following key result enables us to show that $\mathbf{B}M$ is exact only if M is a group.

Theorem 20.4 *Let M be a finite monoid. Then for each $r \in M$ there is a matrix $A \in (\mathbf{B}M)^{2\times 2}$ satisfying*

(i) $[\,M\ \emptyset\,] \in \operatorname{Ker}^2 \operatorname{Row}(A)$*; and*

(ii) $[\,M\ \emptyset\,] \in \operatorname{Row}(A)$ *if and only if* r *has a left inverse in* M*.*

Proof Let $r \in M$. Then since M is finite there is some $n \in \mathbf{N}$ for which r^n is idempotent (see Howie [41, section 1.2]). Now take

$$A = \begin{bmatrix} \{r\} & \emptyset \\ \{1\} & \{r^n\} \end{bmatrix} \in (\mathbf{B}M)^{2\times 2}. \tag{20.5}$$

(i). By Definition 8.3, to conclude that $[\,M\ \emptyset\,] \in \operatorname{Ker}^2 \operatorname{Row}(A)$ we must show that $\operatorname{Ker} \operatorname{Row}(A) \subseteq \operatorname{Ker}([\,M\ \emptyset\,])$, so let $V, V', W, W' \in \mathbf{B}M$ and suppose that

$$A \begin{bmatrix} V \\ W \end{bmatrix} = A \begin{bmatrix} V' \\ W' \end{bmatrix}. \tag{20.6}$$

We then need to check that

$$\begin{bmatrix} M & \emptyset \end{bmatrix} \begin{bmatrix} V \\ W \end{bmatrix} = \begin{bmatrix} M & \emptyset \end{bmatrix} \begin{bmatrix} V' \\ W' \end{bmatrix}. \tag{20.7}$$

That is, we need to check that

$$MV = MV \cup \emptyset W = MV' \cup \emptyset W' = MV', \tag{20.8}$$

and thus by symmetry it is sufficient to check that $MV \subseteq MV'$. If $V = \emptyset$ then we are done, so suppose that $V \neq \emptyset$ and let $qs \in MV$. Now if $s \in V'$ then we are done because $qs \in MV'$, so suppose further that $s \notin V'$.

Combining (20.6) with the definition of A gives

$$\begin{bmatrix} rV \\ V \cup r^n W \end{bmatrix} = \begin{bmatrix} rV' \\ V' \cup r^n W' \end{bmatrix} \tag{20.9}$$

In particular, we have $rV = rV'$, and thus we can write $rs = rs'$ for some $s' \in V'$ because $s \in V$. Our assumption that $s \in V$ also means that $s \in V \cup r^n W$, so (20.9) gives $s \in V' \cup r^n W'$. Therefore $s \in r^n W'$ because we are assuming that $s \notin V'$, and as such $s = r^n t'$ for some $t' \in W'$. Since r^n is idempotent we then have

$$s = r^n t' = r^n r^n t' = r^n s = r^{n-1} rs = r^{n-1} rs', \tag{20.10}$$

where $s' \in V'$, and as such $qs \in MV'$. Hence $MV \subseteq MV'$, as required.

(ii). By (20.5), $[\, M \;\; \emptyset \,] \in \mathrm{Row}(A)$ if and only if there are $V, W \in \mathbf{B}M$ with

$$\begin{bmatrix} M & \emptyset \end{bmatrix} = \begin{bmatrix} V & W \end{bmatrix} A = \begin{bmatrix} Vr \cup W & Wr^n \end{bmatrix}. \tag{20.11}$$

This happens if and only if there is some $V \in \mathbf{B}M$ with $Vr = M$ (because $Wr^n = \emptyset$ if and only if $W = \emptyset$), and thus $[\, M \;\; \emptyset \,] \in \mathrm{Row}(A)$ if and only if there is some $s \in M$ with $sr = 1$. Hence $[\, M \;\; \emptyset \,] \in \mathrm{Row}(A)$ if and only if r has a left inverse in M. □

Corollary 20.5 *Let M be a finite monoid. Then $\mathbf{B}M$ is exact if and only if M is a group.*

Proof ($\Rightarrow$). Suppose that $\mathbf{B}M$ is exact and let $r \in M$. Then by Theorem 20.4 there is some $A \in (\mathbf{B}M)^{2\times 2}$ with $[\, M \;\; \emptyset \,] \in \mathrm{Ker}^2 \mathrm{Row}(A)$, and with $[\, M \;\; \emptyset \,] \in \mathrm{Row}(A)$ if and only if r has a left inverse in M. Now, since $\mathbf{B}M$ is exact, Proposition 9.10 gives $\mathrm{Ker}^2 \mathrm{Row}(A) = \mathrm{Row}(A)$, and thus we have $[\, M \;\; \emptyset \,] \in \mathrm{Row}(A)$. Therefore r has a left inverse in M. Hence M is a group.

($\Leftarrow$). Suppose that M is a group. Then $\mathbf{B}M$ is exact by Corollary 20.2. □

That groups are the only finite monoids M for which $\mathbf{B}M$ is exact is a non-trivial outcome, and to prove it we really do need to use a result such as Theorem 20.4. In other words, to show that exactness of $\mathbf{B}M$ implies that M is group it is necessary to use the fact that (right) exactness tells us about row spaces of matrices over $\mathbf{B}M$, not just finitely generated ideals of $\mathbf{B}M$. The reason we have to use the full power of exactness is that there exist finite non-group monoids M with $\mathrm{Ker}^2(X) = X$ for all (finitely generated) left ideals $X \subseteq \mathbf{B}M$, and so at this level $\mathbf{B}M$ has precisely the exactness-like properties it would have if M were a group. The following example illustrates this obstruction in the case of a particular such a monoid.

Example 20.6 The two-element monoid $M = \{1, r\}$ with multiplication given by

$$\begin{array}{c|cc} \cdot & 1 & r \\ \hline 1 & 1 & r \\ r & r & r \end{array} \tag{20.12}$$

is not a group, so by Corollary 20.5 $\mathbf{B}M$ is not exact.[1] Direct computation reveals

[1] In section 16 we referred to M as the 'Boolean monoid' and the 'two-element semilattice'.

that the ideals of $\mathbf{B}M$ are

$$\{\emptyset\} = (\mathbf{B}M)\emptyset, \tag{20.13}$$

$$\{\emptyset, \{r\}\} = (\mathbf{B}M)\{r\}, \tag{20.14}$$

$$\{\emptyset, \{r\}, \{1, r\}\} = (\mathbf{B}M)\{1, r\} \tag{20.15}$$

and

$$\{\emptyset, \{1\}, \{r\}, \{1, r\}\} = (\mathbf{B}M)\{1\}, \tag{20.16}$$

and thus each ideal of $\mathbf{B}M$ is principal. Furthermore, each ideal of $\mathbf{B}M$ is generated by an idempotent because we happen to have $VV = V$ for all $V \in \mathbf{B}M$.

Now, if S is a semiring and $a \in S$ is idempotent then the left ideal of S generated by a satisfies $\mathrm{Ker}^2(Sa) = Sa$. To prove this, let $b \in \mathrm{Ker}^2(Sa)$ and observe that $\left(1_{\{a,b\}}, a\right) \in \mathrm{Ker}(Sa)$ because $a1_{\{a,b\}} = a = aa$. Since $b \in \mathrm{Ker}^2(Sa)$, Definition 8.3 then gives $\left(1_{\{a,b\}}, a\right) \in \mathrm{Ker}(b)$, and as such $b = b1_{\{a,b\}} = ba$. Hence $\mathrm{Ker}^2(Sa) \subseteq Sa$, as required.

This result tells us that $\mathrm{Ker}^2(X) = X$ for all ideals $X \subseteq \mathbf{B}M$ because, as we observed above, each ideal of $\mathbf{B}M$ is generated by an idempotent. If M were a group we would instead obtain $\mathrm{Ker}^2(X) = X$ for all ideals $X \subseteq \mathbf{B}M$ from exactness of $\mathbf{B}M$, and so in this respect M is indistinguishable from a group.

Bibliography

[1] M. Akian, R. Bapat, and S. Gaubert. Max-plus algebra. In L. Hogben, editor, *Handbook of Linear Algebra*, Discrete Mathematics and Its Applications, chapter 25. Chapman and Hall, Boca Raton, 2006.

[2] C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis*. Springer, Berlin, third edition, 2006.

[3] S. Awodey. *Category Theory*, volume 52 of *Oxford Logic Guides*. Oxford University Press, Oxford, second edition, 2010.

[4] F. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity: An Algebra for Discrete Event Systems*. John Wiley & Sons, New York, 1992.

[5] R. Baer. Abelian groups that are direct summands of every containing abelian group. *Bull. Amer. Math. Soc.*, 46:800–806, 1940.

[6] M. Barr. **-Autonomous Categories*, volume 752 of *Lecture Notes in Mathematics*. Springer, Berlin, 1979.

[7] J. Berstel and C. Reutenauer. *Noncommutative Rational Series with Applications*. Cambridge University Press, Cambridge, 2011.

[8] G. Birkhoff. *Lattice Theory*. American Mathematical Society, New York, second edition, 1948.

[9] T. S. Blyth. *Lattices and Ordered Algebraic Structures*. Springer, London, 2005.

[10] T. S. Blyth and M. F. Janowitz. *Residuation Theory*, volume 102 of *International Series of Monographs in Pure and Applied Mathematics*. Pergamon Press, Oxford, 1972.

[11] G. Boole. *The Mathematical Analysis of Logic*. Macmillan, Barclay and Macmillan, Cambridge, 1847.

[12] J. W. Brewer, J. W. Bunce, and F. S. Van Vleck. *Linear Systems over Commutative Rings*, volume 104 of *Lecture Notes in Pure and Applied Mathematics*. Marcel Dekker, New York, 1986.

[13] R. Bělohlávek. Concept lattices and order in fuzzy logic. *Ann. Pure Appl. Logic*, 128(1–3):277–298, 2004.

[14] R. Bělohlávek and J. Konečný. Row and column spaces of matrices over residuated lattices. *Fund. Inform.*, 115(4):279–295, 2012.

[15] S. Bulman-Fleming and M. Mahmoudi. The category of S-posets. *Semigroup Forum*, 71(3):443–461, 2005.

[16] P. Butkovič. *Max-linear Systems: Theory and Algorithms*. Springer, London, 2010.

[17] A. H. Clifford and G. B. Preston. *The Algebraic Theory of Semigroups, Volume I*, volume 7 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, 1961.

[18] G. Cohen, S. Gaubert, and J.-P. Quadrat. Linear projectors in the max-plus algebra. In *Proceedings of the 5th IEEE Mediterranean Conference on Control and Systems*. IEEE Press, New York, 1997.

[19] G. Cohen, S. Gaubert, and J.-P. Quadrat. Hahn-Banach separation theorem for max-plus semimodules. In J. L. Menaldi, E. Rofman, and A. Sulem, editors, *Optimal Control and Partial Differential Equations*, pages 325–334. IOS Press, Amsterdam, 2001.

[20] G. Cohen, S. Gaubert, and J.-P. Quadrat. Duality and separation theorems in idempotent semimodules. *Linear Algebra Appl.*, 379:395–422, 2004.

[21] P. M. Cohn. *Algebra, Volume 1*. John Wiley & Sons, Chichester, 1974.

[22] R. Cuninghame-Green. *Minimax Algebra*, volume 166 of *Lecture Notes in Economics and Mathematical Systems*. Springer, Berlin, 1979.

[23] F. d'Alessandro and E. Pasku. A combinatorial property for semigroups of matrices. *Semigroup Forum*, 67(1):22–30, 2003.

[24] M. Develin and B. Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27, 2004.

[25] M. Develin, F. Santos, and B. Sturmfels. On the rank of a tropical matrix. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*, volume 52 of *Mathematical Sciences Research Institute Publications*, pages 213–242. Cambridge University Press, Cambridge, 2005.

[26] L. E. Dickson. Definitions of a group and a field by independent postulates. *Trans. Amer. Math. Soc.*, 6(2):198–204, 1905.

[27] N. Galatos, P. Jipsen, T. Kowalski, and H. Ono. *Residuated Lattices: An Algebraic Glimpse at Substructural Logics*, volume 151 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, Amsterdam, 2007.

[28] B. Ganter and R. Wille. *Formal Concept Analysis: Mathematical Foundations*. Springer, Berlin, 1999.

[29] S. Gaubert and R. D. Katz. Reachability problems for products of matrices in semirings. *Internat. J. Algebra Comput.*, 16(3):603–627, 2006.

[30] S. Gaubert and M. Sharify. Tropical scaling of polynomial matrices. In R. Bru and S. Romero-Vivó, editors, *Positive Systems*, volume 389 of *Lecture Notes in Control and Information Sciences*, pages 291–303. Springer, Berlin, 2009.

[31] K. Głazek. *A Guide to the Literature on Semirings and their Applications in Mathematics and Information Sciences*. Kluwer Academic Publishers, Dordrecht, 2002.

[32] J. S. Golan. *Power Algebras over Semirings*. Kluwer Academic Publishers, Dordrecht, 1999.

[33] J. S. Golan. *Semirings and Affine Equations over Them*. Kluwer Academic Publishers, Dordrecht, 2003.

[34] J. S. Golan. *Semirings and their Applications*. Kluwer Academic Publishers, Dordrecht, 2010.

[35] J. A. Green. On the structure of semigroups. *Ann. of Math. (2)*, 54(1):163–172, 1951.

[36] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.

[37] B. Heidergott, G. J. Olsder, and J. van der Woude. *Max Plus at Work*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, 2006.

[38] M. Henriksen. Some remarks on elementary divisor rings. II. *Michigan Math. J.*, 3(2):159–163, 1955.

[39] D. Hilbert. Über den zahlbegriff. *Jahresber. Dtsch. Math.-Ver.*, 8:180–184, 1900.

[40] C. Hollings and M. Kambites. Tropical matrix duality and Green's $\mathcal{D}$ relation. *J. Lond. Math. Soc.*, 86(2):520–538, 2012.

[41] J. M. Howie. *Fundamentals of Semigroup Theory*. Clarendon Press, Oxford, 1995.

[42] E. V. Huntington. A complete set of postulates for the theory of absolute continuous magnitude. *Trans. Amer. Math. Soc.*, 3(2):264–279, 1902.

[43] E. V. Huntington. Complete sets of postulates for the theories of positive integral and positive rational numbers. *Trans. Amer. Math. Soc.*, 3(2):280–284, 1902.

[44] Z. Izhakian and S. W. Margolis. Semigroup identities in the monoid of two-by-two tropical matrices. *Semigroup Forum*, 80(2):191–218, 2010.

[45] N. Jacobson. *Basic Algebra I*. W. H. Freeman and Company, San Francisco, second edition, 1985.

[46] M. Johnson and M. Kambites. Multiplicative structure of 2×2 tropical matrices. *Linear Algebra Appl.*, 435(7):1612–1625, 2010.

[47] M. Johnson and M. Kambites. Green's $\mathcal{J}$-order and the rank of tropical matrices. *J. Pure Appl. Algebra*, 217(2):280–292, 2013.

[48] M. Johnson and M. Kambites. Idempotent tropical matrices and finite metric spaces. *Adv. Geom.*, 14(2):253–276, 2014.

[49] I. Kaplansky. Elementary divisors and modules. *Trans. Amer. Math. Soc.*, 66: 464–491, 1949.

[50] G. M. Kelly. *Basic Concepts of Enriched Category Theory*, volume 64 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1982.

[51] M. Kilp, U. Knauer, and A. V. Mikhalev. *Monoids, Acts and Categories: With Applications to Wreath Products and Graphs*, volume 29 of *de Gruyter Expositions in Mathematics*. Walter de Gruyter, Berlin, 2000.

[52] K. H. Kim. *Boolean Matrix Theory and Applications*. Marcel Dekker, New York, 1982.

[53] S. C. Kleene. Representation of events in nerve nets and finite automata. In C. E. Shannon and J. McCarthy, editors, *Automata Studies*, volume 34 of *Annals of Mathematics Studies*, pages 3–41. Princeton University Press, Princeton, 1956.

[54] W. Krull. Axiomatische begründung der algemeinen idealtheorie. *Sitzungsberichte der physikalischmedizinischen Societät zu Erlangen*, 56:47–63, 1924.

[55] W. Kuich and A. Salomaa. *Semirings, Automata, Languages*, volume 5 of *EATCS Monographs on Theoretical Computer Science*. Springer, Berlin, 1986.

[56] T. Y. Lam. *Lectures on Modules and Rings*, volume 189 of *Graduate Texts in Mathematics*. Springer, New York, 1999.

[57] T. Y. Lam. *A First Course in Noncommutative Rings*, volume 131 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2001.

[58] J. Lambek. *Lectures on Rings and Modules*. Blaisdell, Waltham, Mass., 1966.

[59] F. W. Lawvere. Metric spaces, generalized logic, and closed categories. *Rend. Sem. Mat. Fis. Milano*, 43(1):135–166, 1973.

[60] P. D. Lax. *Functional Analysis*. John Wiley & Sons, New York, 2002.

[61] G. Litvinov. Maslov dequantization, idempotent and tropical mathematics: A brief introduction. *J. Math. Sci.*, 140(3):426–444, 2007.

[62] G. L. Litvinov, V. P. Maslov, and G. B. Shpiz. Idempotent functional analysis: An algebraic approach. *Math. Notes*, 69(5):696–729, 2001.

[63] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*, volume 89 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, Berlin, 1977.

[64] D. Maclagan and B. Sturmfels. *Introduction to Tropical Geometry*, volume 161 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 2015.

[65] B. H. Maddox. Absolutely pure modules. *Proc. Amer. Math. Soc.*, 18:155–158, 1967.

[66] W. W. McGovern. Bézout rings with almost stable range 1. *J. Pure Appl. Algebra*, 212(2):340–348, 2008.

[67] P.-A. Melliès. Categorial semantics of linear logic. In *Interactive Models of Computation and Program Behavior*, volume 27 of *Panoramas et Synthèses*. Société Mathématique de France, Marseilles, 2009.

[68] G. Mikhalkin. Enumerative tropical algebraic geometry in $\mathbb{R}^2$. *J. Amer. Math. Soc.*, 18(2):313–377, 2005.

[69] W. K. Nicholson. Lifting idempotents and exchange rings. *Trans. Amer. Math. Soc.*, 229:269–278, 1977.

[70] W. K. Nicholson and M. F. Yousif. *Quasi-Frobenius Rings*, volume 158 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2003.

[71] J.-E. Pin. Tropical semirings. In J. Gunawardena, editor, *Idempotency*, volume 11 of *Publications of the Newton Institute*, pages 50–69. Cambridge University Press, Cambridge, 1998.

[72] S. Roman. *Advanced Linear Algebra*, volume 135 of *Graduate Texts in Mathematics*. Springer, New York, third edition, 2008.

[73] J. J. Rotman. *An Introduction to Homological Algebra*. Academic Press, London, 1979.

[74] M. P. Schützenberger. On the definition of a family of automata. *Inf. Control*, 4(2–3):245–270, 1961.

[75] R. Y. Sharp. *Steps in Commutative Algebra*, volume 51 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, second edition, 2000.

[76] Y. Shitov. Group rings that are exact. *J. Algebra*, 403:179–184, 2014.

[77] H. Simmons. *An Introduction to Category Theory*. Cambridge University Press, Cambridge, 2011.

[78] I. Simon. Recognizable sets with multiplicities in the tropical semiring. In M. P. Chytil, V. Koubek, and L. Janiga, editors, *Mathematical Foundations of Computer Science 1988*, volume 324 of *Lecture Notes in Computer Science*, pages 107–120. Springer, Berlin, 1988.

[79] I. Simon. On semigroups of matrices over the tropical semiring. *Inform. Théor. Appl.*, 28(3–4):277–294, 1994.

[80] M. Sipser. *Introduction to the Theory of Computation*. Cengage Learning, Boston, third edition, 2013.

[81] H. J. S. Smith. On systems of linear indeterminate equations and congruences. *Philos. Trans. R. Soc. Lond.*, 151:293–326, 1861.

[82] H. S. Vandiver. Note on a simple type of algebra in which the cancellation law of addition does not hold. *Bull. Amer. Math. Soc.*, 40(12):914–920, 1934.

[83] H. Wang. Injective hulls of semimodules over additively-idempotent semirings. *Semigroup Forum*, 48:377–379, 1994.

[84] M. Ward and R. P. Dilworth. Residuated lattices. *Trans. Amer. Math. Soc.*, 45:335–354, 1939.

[85] A. N. Whitehead. *A Treatise on Universal Algebra*. Cambridge University Press, Cambridge, 1898.

[86] D. Wilding, M. Johnson, and M. Kambites. Exact rings and semirings. *J. Algebra*, 388:324–337, 2013.

[87] A. M. Wille. *Residuated Structures with Involution*. PhD thesis, Technische Universität Darmstadt, 2006.

[88] R. Wille. Restructuring lattice theory: An approach based on hierarchies of concepts. In I. Rival, editor, *Ordered Sets*, volume 83 of *NATO Advanced Study Institutes Series*, pages 445–470. Springer, Dordrecht, 1982.

Index of terminology